arXiv:1510.02642v4 [cs.LO] 11 Feb 2016

An Instantiation-Based Approach for Solving Quantified Linear Arithmetic

Andrew Reynolds[1], Tim King[2], and Viktor Kuncak[1]

[1] École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[2] Verimag

Abstract. This paper presents a framework to derive instantiation-based decision procedures for satisfiability of quantified formulas in first-order theories, including its correctness, implementation, and evaluation. Using this framework we derive decision procedures for linear real arithmetic (LRA) and linear integer arithmetic (LIA) formulas with one quantifier alternation. Our procedure can be integrated into the solving architecture used by typical SMT solvers. Experimental results on standardized benchmarks from model checking, static analysis, and synthesis show that our implementation of the procedure in the SMT solver CVC4 outperforms existing tools for quantified linear arithmetic.

1 Introduction

Among the biggest challenges in automated reasoning is efficient support for quantifiers in the presence of background theories. Quantifiers enable direct encoding of a number of problems of interest, including synthesis of software fragments from specifications [31, 47, 53], construction of transfer functions for program analysis [37], invariant inference [13, 25], as well as analysis of properties that go beyond safety [9, 10].

The most commonly used complete method for deciding constraints over quantified theories is *quantifier elimination* [27, Section 2.7]. Quantifier elimination algorithms typically solve a more general problem, of transforming arbitrary quantified formula with free variables into a theory-equivalent formula with no quantifiers. However, depending on the particular variant of the language of constraints, performing actual quantifier elimination can have worse complexity than the decision problem [8], in part because it is required to give an answer on any formula, and the smallest formula resulting from quantifier elimination can be very large [61]. When the goal is to decide the satisfiability of quantified constraints, quantifier elimination may be doing unnecessary work. More importantly, procedures based on quantifier elimination often do not handle the underlying ground constraints in the most efficient way. Thus, quantifier elimination tends to be prohibitively expensive in practice. Recent work involving quantifier elimination [11, 38] has been motivated by avoiding worst-case performance by effectively computing an equisatisfiable set of ground formulas in a lazy fashion.

In the broader scope of automated theorem proving, it is often important to reason about formulas involving multiple theories, each of which may or may not support quantifier elimination. In practice, the goal is to obtain a framework for handling quantified formulas that is both complete for formulas belonging to decidable logics, and empirically effective when completeness guarantees are not known. To this end, modern SMT

solvers most commonly use heuristic instantiation-based approaches [17], which are incomplete but work well in practice for undecidable fragments of first-order logic.

Thus, our motivation is to capitalize both on recent advances in specialized techniques for quantified linear arithmetic [11, 12, 29, 41], and recent advances in instantiation-based theorem proving for first-order logic [16, 22, 49]. This paper seeks to bridge the gap between these two lines of research by introducing an approach for establishing the satisfiability of formulas in quantified linear arithmetic based on a new *quantifier instantiation* framework. The use of quantifier instantiation for this task is motivated by the following.

- Procedures based on lazy quantifier instantiation typically establish satisfiability much faster than their theoretical complexity.
- Using quantifier instantiation for decidable fragments enables a uniform integration and composition with existing instantiation-based techniques [16,17,49], which are widely used by modern SMT solvers.
- An important class of synthesis problems can be expressed as quantified formulas with one quantifier alternation. As shown in [47], solutions for these problems can be extracted from an unsatisfiable core of quantifier instantiations.

Related Work Quantifier elimination has been used to, e.g., show decidability and classification of boolean algebras [52, 58], Presburger arithmetic [43], decidability of products [20, 39], [36, Chapter 12], and algebraically closed fields [57]. The original result on decidability of Presburger arithmetic is by Presburger [43]. The space bound for Presburger arithmetic was shown in [21]. The matching lower and upper bounds for Presburger arithmetic were shown in [8], see also [30, Lecture 24]. An analysis parameterized by the number of quantifier alternations is presented in [45]. A mechanically verified quantifier elimination algorithm was developed by Nipkow [40].

An approach for lazy quantifier elimination for linear real arithmetic was developed by Monniaux [38]. Integration of linear quantifier elimination into the solving algorithm used by SMT solvers was developed in [11], though the presented integration is not model driven. A lazy approach for quantifier elimination, which relies on an operation called model-based projection, has been developed in the context of SMT-based model checking [29], and can be used for extracting skolem functions for simulation synthesis [19]. A recent approach for quantified formulas with arbitrary alternations has been developed by Bjorner [12] for several background theories, which is not based on instantiation. The most widely used techniques for quantifier instantiation in SMT were developed in [17], and later in [16, 23], which primarily focused on uninterpreted functions. Our approach for quantified linear arithmetic instantiates quantified formulas based on a lazy stream of candidate models, terminating when either it finds a finite set of instances are unsatisfiable, or discovers that the original formula is satisfiable. Other approaches in this spirit have been used to decide essentially uninterpreted fragment [24], and, more generally, theories having a locality property [3, 28]; these works do not directly apply to quantified linear arithmetic. A recent approach for quantified formulas with one quantifier alternation has been developed in the SMT solver Yices [18], which does not treat linear integer arithmetic. The present paper builds upon our previous work for solving synthesis conjectures using quantifier instantiation in

SMT [47], where an approach for quantified linear arithmetic was described without a specific method for selecting instances and without completeness guarantees. While the present paper focuses on linear arithmetic, where it outperforms existing approaches, we expect the presented framework to be relevant for other quantified theories. Among the examples of further decidable quantified constraints are quantified theories of term algebras [36, Chapter 23], [35, 54] and their extensions [14, 32, 50], feature trees [2, 59], and monadic second-order theories [60].

Contributions This paper makes the following contributions. First, we define a general class of instantiation-based procedures for establishing the satisfiability of quantified formulas in Section 2, and show that these procedures can be used in part as an approach for solving synthesis problems in Section 2.3. We demonstrate instances of the procedure are sound and complete for formulas over linear real arithmetic (LRA) and linear integer arithmetic (LIA) with one quantifier alternation in Sections 3 and 4, two quantified fragments for which many current SMT solvers do not have efficient support for. We show how our procedure can be integrated into the solving architecture used by SMT solvers in Section 5. To our knowledge, our approach is the first complete algorithm for quantified linear arithmetic with one alternation that is based purely on quantifier instantiation, which has the advantage of being composable with existing techniques and whose soundness is straightforward to verify. In Section 6, we demonstrate an implementation of the procedures for LIA and LRA in the SMT solver CVC4, which in addition to having the aforementioned advantages, outperforms state-of-the-art SMT solvers and theorem provers for quantified linear arithmetic benchmarks.

1.1 Preliminaries

We consider formulas in multi-sorted first-order logic. A *signature* Σ consists of a countable set of sort symbols and a set of function symbols. Given a signature Σ, well-sorted terms, atoms, literals, and formulas are defined as usual, and referred to respectively as *Σ-terms*. We denote by $FV(t)$ the set of free variables occurring in the term t, and extend this notion to formulas. A Σ-term or formula is *ground* if it has no free variables. A term written $t[\boldsymbol{k}]$ denotes a term whose free variables are in $\boldsymbol{k}$.

A *Σ-interpretation* $\mathcal{I}$ maps

- each set sort symbol $\sigma \in \Sigma$ to a non-empty set $\sigma^{\mathcal{I}}$, the *domain* of σ in $\mathcal{I}$, and
- each function $f \in \Sigma$ of sort $\sigma_1 \times \ldots \times \sigma_n \to \sigma$ to a total function $f^{\mathcal{I}}$ of sort $\sigma_1^{\mathcal{I}} \times \ldots \times \sigma_n^{\mathcal{I}} \to \sigma^{\mathcal{I}}$ where $n > 0$, and to an element of $\sigma^{\mathcal{I}}$ when $n = 0$.

We write $t^{\mathcal{I}}$ to denote the interpretation of t in $\mathcal{I}$, defined inductively as usual. A satisfiability relation between Σ-interpretations and Σ-formulas, written $\mathcal{I} \models \varphi$, is also defined inductively as usual. In particular, we assume that $\mathcal{I} \models \neg\varphi$ if and only if it is not the case that $\mathcal{I} \models \varphi$. We say that $\mathcal{I}$ is *a model of* φ if $\mathcal{I}$ satisfies φ. Formulas φ_1 and φ_2 are *equivalent (up to $\boldsymbol{k}$)* if they are satisfied by the same set of models (when restricted to the interpretation of variables $\boldsymbol{k}$).

A *theory* is a pair $T = (\Sigma, \mathbf{I})$ where Σ is a signature and $\mathbf{I}$ is a non-empty set of Σ-interpretations, the *models* of T. We assume Σ contains the equality predicate, which we denote by $\approx$. Let $[\![\varphi]\!]_T$ denote the set of T-models of φ. Observe that $[\![\neg\varphi]\!]_T =$

$\mathbf{I} \setminus [\![\varphi]\!]_T$. A Σ-formula $\varphi[\boldsymbol{x}]$ is *T-satisfiable* if it is satisfied by some interpretation in $\mathbf{I}$ (i.e. $[\![\varphi]\!]_T \neq \emptyset$). Dually, a Σ-formula $\varphi[\boldsymbol{x}]$ is *T-unsatisfiable* if it is satisfied by no interpretation in $\mathbf{I}$ (i.e. $[\![\varphi]\!]_T = \emptyset$). A formula φ is *T-valid* if every model of T is a model of φ (i.e., $[\![\varphi]\!]_T = \mathbf{I}$). Given a fragment $\mathbf{L}$ of the language of Σ-formulas, a Σ-theory T is *satisfaction complete with respect to* $\mathbf{L}$ if every closed T-satisfiable formula of $\mathbf{L}$ is T-valid. In terms of set of models, satisfaction completeness means that $[\![\varphi]\!]_T \neq \emptyset$ implies $[\![\varphi]\!]_T = \mathbf{I}$, or, in other words, for every $F \in \mathbf{L}$ exactly one of the following two cases hold: $[\![\varphi]\!]_T = \emptyset$, or $[\![\varphi]\!]_T = \mathbf{I}$. If additionally $\mathbf{L}$ is closed under negation, then, for every $\varphi \in \mathbf{L}$, either φ or $\neg F$ is unsatisfiable.

A set Γ of formulas *T-entails* a Σ-formula φ, written $\Gamma \models_T \varphi$, if every model of T that satisfies all formulas in Γ satisfies φ as well. A set of literals M *propositionally entails* a formula φ, written $M \models_p \varphi$, if M entails φ when considering all atomic formulas in $M \cup \varphi$ as propositional variables; such entailment is one of from propositional logic and is independent of the theory.

We write RA (resp. IA) to denote the theory of real (resp. integer) arithmetic. Its signature consists of the sort Real (resp. Int), the binary predicate symbols $>$ and $<$, functions $+$ and $\cdot$ denoting addition and multiplication, and the constants of its sort interpreted as usual. We write $t \leq s$ as shorthand for $\neg(t > s)$, and $t \geq s$ as shorthand for $\neg(t < s)$. We write LRA (resp. LIA) to denote the language of linear real (resp. integer) arithmetic formulas, that is, whose literals are of the form $(\neg)(c_1 \cdot x_1 + \ldots + c_n \cdot x_n \bowtie c)$ where $c_1, \ldots, c_n, c$ and $x_1, \ldots, x_n$ are non-zero constants and distinct variables of sort Real (resp. Int) respectively, and $\bowtie$ is one of $>$, $<$, or $\approx$. For each literal of this form, there exists an equivalent literal that is in *solved form with respect to* x_i for each $i = 1, \ldots, n$. That is, an LRA-literal is in solved form with respect to x if it is of the form $(\neg)(x \bowtie t)$, where $x \notin FV(t)$. Similarly, an LIA-literal is in solved form with respect to x if it is of the form $(\neg)(c \cdot x \bowtie t)$, where $x \notin FV(t)$ and c is an integer constant greater than zero. For integer constants c_1 and c_2 and non-zero constant c, we write $c_1 \equiv_c c_2$ to denote that c_1 and c_2 are congruent modulo c, that is $(c_1 \bmod c) = (c_2 \bmod c)$, and we write $c \mid c_1$ if c divides c_1.

2 Quantifier Instantiation for Theories

In this section, we assume a fixed theory T and a language $\mathbf{L}$ that is closed under negation and such that the satisfiability of finite sets of $\mathbf{L}$ formulas modulo T is decidable. We present a procedure for checking satisfiability of formulas in the language $\mathcal{Q}(\mathbf{L}) = \{\forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}] \mid \varphi[\boldsymbol{a}, \boldsymbol{x}] \in \mathbf{L}\}$.

2.1 An Instantiation Procedure and Its Soundness

Figure 1 presents an instantiation-based approach for determining the satisfiability of a T-formulas $\exists \boldsymbol{a} \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$, where $\varphi[\boldsymbol{a}, \boldsymbol{x}]$ belongs to $\mathbf{L}$. The procedure introduces a tuple of distinct fresh constants $\mathbf{k}$ of the same sort as $\boldsymbol{a}$, and $\mathbf{e}$ of the same sort as $\boldsymbol{x}$. It maintains a set of formulas Γ, initially empty, and terminates when either Γ or $\Gamma \cup \{\neg\varphi[\mathbf{k}, \mathbf{e}]\}$ is T-unsatisfiable. On each iteration, the procedure invokes the subprocedure $\mathcal{S}$ (over which the procedure is parameterized), which returns a tuple of terms

$\mathcal{P}_{\mathcal{S}}(\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}])$:

Let $\Gamma := \emptyset$ and $\mathbf{k}$, $\mathbf{e}$ be tuples of fresh constants of the same type as $\boldsymbol{a}$, $\boldsymbol{x}$.
Repeat
  If Γ is T-unsatisfiable, then return "unsat".
  If $\Gamma' = \Gamma \cup \{\neg\varphi[\mathbf{k}, \mathbf{e}]\}$ is T-unsatisfiable, then return "sat".
  Otherwise,
    Let $\mathcal{I}$ be a model of T and Γ' and let $\boldsymbol{t}[\mathbf{k}] = \mathcal{S}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}])$.
    $\Gamma := \Gamma \cup \{\varphi[\mathbf{k}, \boldsymbol{t}[\mathbf{k}]]{\downarrow}\}$.

Fig. 1. An instantiation-based approach $\mathcal{P}_{\mathcal{S}}$ for determining the T-satisfiability of $\exists \boldsymbol{a}\, \forall x\, \varphi[\boldsymbol{a}, x]$ parameterized by selection function $\mathcal{S}$.

$\boldsymbol{t}[\mathbf{k}]$ whose free variables are a subset of $\mathbf{k}$. We then add to Γ the formula $\varphi[\mathbf{k}, \boldsymbol{t}[\mathbf{k}]]{\downarrow}$, a formula equivalent to $\varphi[\mathbf{k}, \boldsymbol{t}[\mathbf{k}]]$ up to $\mathbf{k}$[3]. We call $\mathcal{S}$ the *selection function* of $\mathcal{P}_{\mathcal{S}}$.

The intution of the algorithm is to find a subset of the instances of $\forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ that are either (a) unsatisfiable, and are thus sufficient for showing that $\forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ is unsatisfiable, or (b) satifiable and entail $\forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$. The algorithm recognizes the latter case by checking the satisfiability of $\Gamma \cup \neg\varphi[\boldsymbol{k}, \boldsymbol{e}]$ on each iteration of its main loop. In either case, the algorithm may terminate before enumerating all instances of $\forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$. In practice, we have found the algorithm often terminates after enumerating only a small number of instances for benchmarks that occur in practice.

Definition 1 *A selection function (for* $\mathbf{L}$*) takes as arguments an interpretation* $\mathcal{I}$*, a set of formulas* Γ*, and a formula* $\neg\varphi[\mathbf{k}, \mathbf{e}]$ *in* $\mathbf{L}$*, where* $\mathcal{I} \models \Gamma \cup \neg\varphi[\mathbf{k}, \mathbf{e}]$*, and returns a tuple of terms* $\boldsymbol{t}[\mathbf{k}]$ *such that* $\varphi[\mathbf{k}, \boldsymbol{t}[\mathbf{k}]]{\downarrow}$ *is also in* $\mathbf{L}$*.*

Note that a selection function is only defined if $\mathcal{I}$ is a model for T, Γ and $\neg\varphi[\mathbf{k}, \mathbf{e}]$. We first show that the procedure always returns correct results, regardless of the behavior of the selection function, leaving the termination question for the next subsection. Detailed proofs of the all claims in this paper can be found in the extended version of this report [48].

Lemma 1 *If* $\mathcal{P}_{\mathcal{S}}$ *terminates with "unsat", then* $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ *is* T*-unsatisfiable.*

Proof: In this case, there exists a set Γ that is equivalent to $\{\varphi[\mathbf{k}, \boldsymbol{t}_1], \ldots, \varphi[\mathbf{k}, \boldsymbol{t}_p]\}$ and is T-unsatisfiable where $\mathbf{k}$ are distinct fresh constants. Thus, $\forall \boldsymbol{x}\, \varphi[\mathbf{k}, \boldsymbol{x}]$ is T-unsatisfiable. Since $\mathbf{k}$ are distinct and fresh, we conclude that $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ is T-unsatisfiable. ■

Lemma 2 *If* $\mathcal{P}_{\mathcal{S}}$ *terminates with "sat", then* $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ *is* T*-satisfiable.*

Proof: In this case, there exists a set Γ equivalent to $\{\varphi[\mathbf{k}, \boldsymbol{t}_1], \ldots, \varphi[\mathbf{k}, \boldsymbol{t}_p]\}$ that is T-satisfiable, where $\mathbf{k}$ are distinct fresh constants, and Γ' equivalent to $\Gamma \cup \{\neg\varphi[\mathbf{k}, \mathbf{e}]\}$ that

[3] We further comment on examples of operators $\downarrow$ in Sections 3 and 4.

is T-unsatisfiable. The variables $\mathbf{e}$ do not occur in Γ, since the only formulas added to it are of the form $\varphi[\mathbf{k}, \boldsymbol{t}[\mathbf{k}]]\downarrow$. Thus, we have that $\Gamma \cup \{\exists \boldsymbol{x}\, \neg\varphi[\mathbf{k}, \boldsymbol{x}]\}$ is T-unsatisfiable. Let $\mathcal{I}$ be a model of Γ. Since $\mathcal{I}$ is not a model of Γ', it must be the case that $\mathcal{I} \not\models \exists \boldsymbol{x}\, \neg\varphi[\mathbf{k}, \boldsymbol{x}]$, and hence $\mathcal{I}$ is a model for $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$. ■

2.2 Termination of the Instantiation Procedure

The following properties of selection functions will be of interest.

Definition 2 (Finite) *A selection function $\mathcal{S}$ is* finite *for $\varphi[\mathbf{k}, \mathbf{e}]$ if there exists a finite set $\mathcal{S}^*(\varphi[\mathbf{k}, \mathbf{e}])$ such that $\mathcal{S}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}]) \in \mathcal{S}^*(\varphi[\mathbf{k}, \mathbf{e}])$ for all $\mathcal{I}$, Γ.*

Definition 3 (Monotonic) *A selection function $\mathcal{S}$ is* monotonic *for $\varphi[\mathbf{k}, \mathbf{e}]$ if whenever $\Gamma \models \varphi[\mathbf{k}, \boldsymbol{t}]$, we have that $\mathcal{S}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}]) \neq \boldsymbol{t}$.*

Observe that, if $\mathcal{S}$ is a monotonic selection function, then for any finite list of terms $\boldsymbol{t}_1, \ldots \boldsymbol{t}_n$ we have $\mathcal{S}(\mathcal{I}, \{\varphi[\mathbf{k}, \boldsymbol{t_1}]\downarrow, \ldots, \varphi[\mathbf{k}, \boldsymbol{t_n}]\downarrow\}, \neg\varphi[\mathbf{k}, \mathbf{e}]) \notin \{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_n\}$.

Definition 4 (Model-Preserving) *A selection function $\mathcal{S}$ is* model-preserving *for $\varphi[\mathbf{k}, \mathbf{e}]$ if whenever $\mathcal{S}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}]) = \boldsymbol{t}$, we have that $\mathcal{I} \models \neg\varphi[\mathbf{k}, \boldsymbol{t}]$.*

Lemma 3 *A selection function that is model-preserving for $\varphi[\mathbf{k}, \mathbf{e}]$ is also monotonic for $\varphi[\mathbf{k}, \mathbf{e}]$.*

Proof: Assume that $\mathcal{S}$ is model-preserving for $\varphi[\mathbf{k}, \mathbf{e}]$ and that $\mathcal{S}(\mathcal{I}, \Gamma \cup \{\varphi[\mathbf{k}, \boldsymbol{t}]\}, \neg\varphi[\mathbf{k}, \mathbf{e}]) = \boldsymbol{s}$. By definition of selection function, we have that $\mathcal{I} \models \varphi[\mathbf{k}, \boldsymbol{t}]$. By definition of model-preserving, we have that $\mathcal{I} \models \neg\varphi[\mathbf{k}, \boldsymbol{s}]$. Thus, $\boldsymbol{s} \neq \boldsymbol{t}$ and $\mathcal{S}$ is monotonic for $\varphi[\mathbf{k}, \mathbf{e}]$. ■

Theorem 1 *If $\mathcal{S}$ is finite and monotonic for $\varphi[\mathbf{k}, \mathbf{e}]$ in $\mathbf{L}$, then $\mathcal{P}_\mathcal{S}$ is a (terminating) decision procedure for the T-satisfiability of $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$.*

Proof: Given a monotonic and finite $\mathcal{S}$, the procedure $\mathcal{P}_\mathcal{S}$ can only execute a finite number of iterations. Assuming a decision procedure for determining the T-satisfiability of T-formulas in $\mathbf{L}$, $\mathcal{P}_\mathcal{S}(\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}])$ must terminate. By Lemmas 1 and 2, $\mathcal{P}_\mathcal{S}$ is a decision procedure for the T-satisfiability of $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$. ■

In this paper we will identify selection functions $\mathcal{S}$ that are finite and monotonic for all $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ residing in fragments $\mathbf{L}$. The fragments we consider are satisfaction complete. We consider satisfaction completeness to be a good guiding principle when choosing candidate logical fragments to which our method can be applied successfully.

2.3 Connection to Synthesis

The connection between quantifier elimination and synthesis has been shown fruitful in previous work [31]; it is one of our motivations for further improving quantified reasoning modulo theories. The procedure mentioned in this section can be used to

synthesize functions from certain classes of specifications. Consider (second-order) T-formulas of the form:

$$\exists \boldsymbol{f}\ \forall \boldsymbol{x}\ \varphi[\boldsymbol{f}, \boldsymbol{x}] \tag{1}$$

of the form $\exists \boldsymbol{f}\ \forall \boldsymbol{x}\ \varphi[\boldsymbol{f}, \boldsymbol{x}]$, where φ is a quantifier-free formula, $\boldsymbol{x} = (x_1, \ldots, x_n)$ is a tuple of variables of sort σ_i for $i = 1, \ldots, n$, and $\boldsymbol{f} = (f_1, \ldots, f_m)$ is a tuple of functions of sort $\sigma_1 \times \ldots \times \sigma_n \to \tau_j$ for $j = 1, \ldots, m$. We call such formulas *synthesis conjectures*. A synthesis conjecture is *single invocation* (over $\mathbf{L}$) if it is equivalent to:

$$\exists \boldsymbol{f}\ \forall \boldsymbol{x}\ \psi[\boldsymbol{x}, \boldsymbol{f}(\boldsymbol{x})] \tag{2}$$

where $\psi[\boldsymbol{x}, \boldsymbol{y}] \in \mathbf{L}$. That is, functions from $\boldsymbol{f}$ are applied to the tuple $\boldsymbol{x}$ only. The formula (2) is equivalent to the (first-order) formula $\forall \boldsymbol{x}\ \exists \boldsymbol{y}\ \psi[\boldsymbol{x}, \boldsymbol{y}]$, whose negation

$$\exists \boldsymbol{x}\ \forall \boldsymbol{y}\ \neg\psi[\boldsymbol{x}, \boldsymbol{y}] \tag{3}$$

is suitable as an input to Figure 1. As observed in [47], solutions for single invocation synthesis conjectures can be extracted from an unsatisfiable core of instantiations when proving the unsatisfiability of (3). In particular, let k be a set of distinct fresh variables of the same sort as $\boldsymbol{x}$, and say the set $\{\neg\psi[\mathsf{k}, \boldsymbol{t}_1[\mathsf{k}]]\downarrow, \ldots, \neg\psi[\mathsf{k}, \boldsymbol{t}_p[\mathsf{k}]]\downarrow\}$ is T-unsatisfiable where $\boldsymbol{t}_i = (t_i^1[\mathsf{k}], \ldots, t_i^m[\mathsf{k}])$ for $i = 1, \ldots, p$. Then:

$$1 \leq j \leq m : f_j = \lambda \boldsymbol{x}.\, \mathsf{ite}(\psi[\boldsymbol{x}, t_p^j[\boldsymbol{x}]], t_p^j[\boldsymbol{x}], (\cdots \mathsf{ite}(\psi[\boldsymbol{x}, t_2^j[\boldsymbol{x}]], t_2^j[\boldsymbol{x}], t_1^j[\boldsymbol{x}])\cdots)) \tag{4}$$

is a solution for $\boldsymbol{f}$ in (2). We use the instantiation-based procedure in Figure 1 for discharging (3). In contrast to prior work [47], we here devise selection functions $\mathcal{S}$ for $\mathbf{L}$ that are finite and monotonic, obtaining a sound and complete method for synthesizing tuples of functions whose specification is a single invocation synthesis conjecture over $\mathbf{L}$. In the following sections, we show such selection functions both for linear real arithmetic (LRA) and linear real arithmetic (LIA).

It is important to note that the solution (4) does not necessarily belong to the language $\mathbf{L}$, since there is no restriction on the selection functions for $\mathbf{L}$ that restricts its return value $\boldsymbol{t}$ to terms in $\mathbf{L}$. For example, in our approach for linear real arithmetic, $\boldsymbol{t}$ may contain a free distinguished constant δ representing an infinitesimal positive value, and in our approach for linear integer arithmetic, $\boldsymbol{t}$ may contain integer division. Additional steps may be necessary in practice for making (4) a computable function.

2.4 Illustration: Instantiation for a Simple Fragment of LRA

We first present a selection function for a restricted class $\mathbf{L}$ of LRA-formulas $\exists \boldsymbol{a}\, \forall x, \varphi[\boldsymbol{a}, x]$, namely whose universal quantifier is over single variable x of sort Real, $\boldsymbol{a}$ are variables of sort Real, and whose (skolemized) body $\varphi[\mathsf{k}, \mathsf{e}]$ is of the form:

$$(\mathsf{e} < \ell_1 \vee \ldots \vee \mathsf{e} < \ell_n \vee \mathsf{e} > u_1 \vee \ldots \vee \mathsf{e} > u_m) \tag{5}$$

where at least one of $\{n, m\}$ is greater than zero, and $\mathsf{e} \notin FV(\ell_1, \ldots, \ell_n, u_1, \ldots, u_m)$. Figure 2 gives a selection function for $\mathcal{S}_{\mathrm{SLRA}}$. It considers the interpretation of terms $\ell_1, \ldots, \ell_n$ and $u_1, \ldots, u_m$ in a model $\mathcal{I}$ of Γ. If $n > 0$, then $\mathcal{S}_{\mathrm{SLRA}}$ returns the ℓ_j whose value is maximal in $\mathcal{I}$. If $n = 0$, then $m > 0$ and $\mathcal{S}_{\mathrm{SLRA}}$ returns the u_i whose value is minimal in $\mathcal{I}$.

$$\mathcal{S}_{SLRA}(\mathcal{I}, \Gamma, \neg\varphi[\mathsf{k}, \mathsf{e}]) \quad \text{where } \varphi[\mathsf{k}, \mathsf{e}] \text{ is } \bigvee_{i=1}^{n} e < l_i \vee \bigvee_{i=1}^{m} e > u_i, \text{ for } n > 0 \text{ or } m > 0$$
$$\text{Return } \begin{cases} \ell_j, & \text{if } n > 0 \text{ and } \mathsf{max}\{\ell_1^{\mathcal{I}}, \ldots, \ell_n^{\mathcal{I}}\} = \ell_j^{\mathcal{I}} \\ u_i, & \text{if } n = 0 \text{ and } \mathsf{min}\{u_1^{\mathcal{I}}, \ldots, u_m^{\mathcal{I}}\} = u_i^{\mathcal{I}} \end{cases}$$

Fig. 2. A selection function $\mathcal{S}_{\mathrm{SLRA}}$ for a simple fragment of LRA.

Lemma 4 *$\mathcal{S}_{\mathrm{SLRA}}$ is finite for $\varphi[\mathsf{k}, \mathsf{e}]$.*

Proof: The terms returned by $\mathcal{S}_{\mathrm{SLRA}}$ are in $\{\ell_1, \ldots, \ell_n\}$ when $n > 0$, and in $\{u_1, \ldots, u_m\}$ when $n = 0$. ■

Lemma 5 *$\mathcal{S}_{\mathrm{SLRA}}$ is monotonic for $\varphi[\mathsf{k}, \mathsf{e}]$.*

Proof: Let $\mathcal{I}$ be a model of LRA and Γ, where $\Gamma \models \varphi[\mathsf{k}, t]$, and assume by contradiction $\mathcal{S}_{\mathrm{SLRA}}$ returns t. Consider the case where $n > 0$ and $t = \ell_i$ for some $i \in \{1, \ldots, n\}$. Since $\mathcal{I} \models \Gamma \cup \neg\varphi[\mathsf{k}, \mathsf{e}]$, it satisfies:

$$\begin{aligned}(\ell_i < \ell_1 \vee \ldots \vee \ell_i < \ell_n \vee \ell_i > u_1 \vee \ldots \vee \ell_i > u_m) \wedge \\ \mathsf{e} \geq \ell_1 \wedge \ldots \wedge \mathsf{e} \geq \ell_n \wedge \mathsf{e} \leq u_1 \wedge \ldots \wedge \mathsf{e} \leq u_m\end{aligned} \tag{6}$$

Assume $\mathcal{I}$ satisfies $\ell_i > u_k$ for some $k \in \{1, \ldots, m\}$. We have $\mathcal{I}$ also satisfies $\mathsf{e} \geq \ell_i$ and $\mathsf{e} \leq u_k$, and thus $\mathsf{e}^{\mathcal{I}} \geq \ell_i^{\mathcal{I}} > u_k^{\mathcal{I}} \geq \mathsf{e}^{\mathcal{I}}$. Thus, $\mathcal{I}$ must satisfy $\ell_i < \ell_{k'}$ for some $k' \in \{1, \ldots, n\}$, and thus $\mathsf{max}\{\ell_1^{\mathcal{I}}, \ldots, \ell_n^{\mathcal{I}}\} \neq \ell_i^{\mathcal{I}}$. Thus, $\mathcal{S}_{\mathrm{SLRA}}(\Gamma) \neq \ell_i$. By symmetrical reasoning when $n = 0, m > 0$ and $t = u_j$ for some $j \in \{1, \ldots, m\}$, we have that $\mathcal{S}_{\mathrm{SLRA}}(\Gamma) \neq u_j$. Thus, $\mathcal{S}_{\mathrm{SLRA}}$ does not return t, and thus is monotonic for $\varphi[\mathsf{k}, \mathsf{e}]$. ■

Example 1. Consider the formula $\forall x\,(x < b \vee x > a)$. The negated skolemized form ($\neg\varphi[\mathsf{k}, \mathsf{e}]$ in Figure 1) of this formula is equivalent to the formula $b \leq e \wedge e \leq a$ where e is a fresh constant. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{SLRA}}}$ on this input is as follows, where Γ is initially $\emptyset$ and on each iteration $\Gamma' = \Gamma \cup \{b \leq e \wedge e \leq a\}$.

#	Γ	Γ'		$t[\mathsf{k}]$	Add to Γ
1	sat	sat	$\mathsf{max}\{b^{\mathcal{I}}\} = b^{\mathcal{I}}$	b	$b < b \vee b > a$
2	sat	unsat			

In step 2, note that $\Gamma' = \{b < b \vee b > a, b \leq e \wedge e \leq a\}$, which is unsat. The run establishes that $\exists ab\,\forall x\,(x > a \vee x < b)$ is LRA-satisfiable. ■

Example 2. Consider the formula $\forall x\,(x < a \vee x < b)$, whose skolemized negation is equivalent to $e \geq a \wedge e \geq b$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{SLRA}}}$ on this input is as follows.

#	Γ	Γ'		$t[\mathsf{k}]$	Add to Γ
1	sat	sat	$\mathsf{max}\{a^{\mathcal{I}}, b^{\mathcal{I}}\} = a^{\mathcal{I}}$	a	$a < a \vee a < b$
2	sat	sat	$\mathsf{max}\{a^{\mathcal{I}}, b^{\mathcal{I}}\} = b^{\mathcal{I}}$	b	$b < a \vee b < b$
3	unsat				

Note that the formula added in step 1, equivalent to $a < b$, ensures that, in the second iteration, $\mathsf{max}\{a^{\mathcal{I}}, b^{\mathcal{I}}\} \neq a^{\mathcal{I}}$. The run establishes that $\exists ab\, \forall x\, (x > a \vee x > b)$ is LRA-unsatisfiable, as expected in a linear order without endpoints. ■

3 Instantiation for Quantifier-Free LRA-Formulas

Consider the case where $\boldsymbol{a}$ and $\boldsymbol{x}$ are vectors of Real variables and $\mathbf{L}$ is the class of formulas $\exists \boldsymbol{a}\, \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$ where $\varphi[\boldsymbol{a}, \boldsymbol{x}]$ is an arbitrary quantifier-free LRA-formula. We assume that equalities are eliminated from φ by the transformation:

$$t \approx 0 \ \rightsquigarrow\ 0 \leq t \wedge 0 \geq t$$

Figure 3 gives a selection function $\mathcal{S}_{\mathrm{LRA}}$ for LRA, which takes an interpretation $\mathcal{I}$, a set of formulas Γ, and the formula $\neg\varphi[\mathbf{k}, \mathbf{e}]$. It invokes the recursive procedure S_R which constructs a term corresponding to each variable in $\mathbf{e}$. Analogous to existing approaches for linear quantifier elimination [33, 40], our approach makes use of non-standard terms for symbolically representing substitutions. In particular, the terms we consider may involve a free distinguished constant δ, representing an infinitesimal positive value. For each variable e_i from $\mathbf{e}$, the procedure S_R invokes the (non-deterministic) subprocedure S_{R0}, which chooses a term corresponding to e_i based on a set of literals M over the atoms of ψ which propositionally entail ψ and are satisfied by $\mathcal{I}$, which we call a *propositionally satisfying assignment* for ψ. We partition M into three sets M_ℓ, M_u and M_c, where M_ℓ contains literals that correspond to lower bounds for e, M_u contains literals that correspond to upper bounds for e, and M_c contains the remaining literals. The sets M_ℓ and M_u are equivalent to sets of literals that are in solved form with respect to e. When M_ℓ contains at least one literal, we may return the lower bound whose value is maximal according to $\mathcal{I}$, and similarly for M_u. If both M_ℓ and M_u are empty, we return the term 0. When S_{R0} returns the term t_i, we apply the substitution $\{\mathsf{e}_i \mapsto t_i\}$ to ψ and $\boldsymbol{t}$, and append t_i to $\boldsymbol{t}$. Terms returned by S_{R0} may involve the constant δ. We define a satisfiability relation between models and formulas involving δ, as well as the max and min function for terms involving δ in the obvious way, such that $(t_1 + c_1 \cdot \delta)^{\mathcal{I}} > (t_2 + c_2 \cdot \delta)^{\mathcal{I}}$ if either $t_1^{\mathcal{I}} > t_2^{\mathcal{I}}$ or both $t_1^{\mathcal{I}} = t_2^{\mathcal{I}}$ and $c_1 > c_2$.

Overall, $\mathcal{S}_{\mathrm{LRA}}$ returns a tuple of terms $\boldsymbol{t}$, after which we add the instance $\varphi[\mathbf{k}, \boldsymbol{t}]{\downarrow}$ to Γ in Figure 1. We assume $\downarrow$ eliminates occurrences of δ by the following transformations, which is inspired by virtual term substitution [33]:

$$\delta < t \rightsquigarrow 0 < t \text{ and } \delta > t \rightsquigarrow 0 \geq t \text{ where } \delta \notin FV(t).$$

Lemma 6 *$\mathcal{S}_{\mathrm{LRA}}$ is finite for $\varphi[\mathbf{k}, \mathbf{e}]$.*

Proof: We first show that only a finite number of terms can be returned by $S_{R0}(\mathcal{I}, (\neg\varphi[\mathbf{k}, \mathbf{e}])\sigma, \mathsf{e}_i)$ for any $\mathcal{I}, \sigma$. Let A be the set of atoms occurring in $\varphi[\mathbf{k}, \mathsf{e}_i]\sigma$. The literals in satisfying assignments of $(\neg\varphi[\mathbf{k}, \mathbf{e}])\sigma$ are over these atoms. Let $\{\mathsf{e}_i < \ell_1, \ldots, \mathsf{e}_i < \ell_n, \mathsf{e}_i > u_1, \ldots, \mathsf{e}_i > u_m\}$ be the set of atoms that are in solved

$\mathcal{S}_{\mathrm{LRA}}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}])$:

Return $S_R(\mathcal{I}, \neg\varphi[\mathbf{k}, \mathbf{e}], \mathbf{e}, ())$

$S_R(\mathcal{I}, \psi, (\mathsf{e}_i, \ldots, \mathsf{e}_n), \boldsymbol{t})$:

If $i > n$, return $\boldsymbol{t}$
Otherwise, let $t_i = S_{R0}(\mathcal{I}, \psi, \mathsf{e}_i)$, $\sigma = \{\mathsf{e}_i \mapsto t_i\}$
Return $S_R(\mathcal{I}, \psi\sigma, (\mathsf{e}_{i+1}, \ldots, \mathsf{e}_n), (\boldsymbol{t}\sigma, t_i))$

$S_{R0}(\mathcal{I}, \psi, \mathsf{e})$:

Let $M = M_\ell \cup M_u \cup M_c$ be such that:
- $\mathcal{I} \models M$ and $M \models_p \psi$,
- $M_\ell \Leftrightarrow \{\mathsf{e} \succ \ell_1, \ldots, \mathsf{e} \succ \ell_n\}$,
- $M_u \Leftrightarrow \{\mathsf{e} \prec u_1, \ldots, \mathsf{e} \prec u_m\}$, and
- $\mathsf{e} \notin FV(\ell_1, \ldots, \ell_n) \cup FV(u_1, \ldots, u_m) \cup FV(M_c)$.

Return one of $\begin{cases} \ell_i + \delta_i^\ell & n > 0, \mathsf{max}\{(\ell_1 + \delta_1^\ell)^\mathcal{I}, \ldots, (\ell_n + \delta_n^\ell)^\mathcal{I}\} = (\ell_i + \delta_i^\ell)^\mathcal{I} \\ u_j - \delta_j^u & m > 0, \mathsf{min}\{(u_1 - \delta_1^u)^\mathcal{I}, \ldots, (u_m - \delta_m^u)^\mathcal{I}\} = (u_j - \delta_j^u)^\mathcal{I} \\ 0 & n = 0 \text{ and } m = 0 \end{cases}$

Fig. 3. A selection function $\mathcal{S}_{\mathrm{LRA}}$ for arbitrary quantifier-free LRA-formula $\varphi[\boldsymbol{a}, x]$. Each $\prec$ is either $<$ or $\leq$; δ_i^ℓ is δ if the i^{th} lower bound for e is strict, and 0 otherwise. Similarly, each $\succ$ is either $>$ or $\geq$; δ_j^u is δ if the j^{th} upper bound for e is strict, and 0 otherwise.

form with respect to e_i that are equivalent to the atoms of A containing e_i, where $\mathsf{e}_i \notin FV(\ell_1, \ldots, \ell_n, u_1, \ldots, u_m)$. The terms returned by $S_{R0}(\mathcal{I}, (\neg\varphi[\mathbf{k}, \mathbf{e}])\sigma, \mathsf{e}_i)$ are in $\{0, \ell_1(+\delta), \ldots, \ell_n(+\delta), u_1(-\delta), \ldots, u_m(-\delta)\}$. Since there are only a finite number of recursive calls to S_R within $\mathcal{S}_{\mathrm{LRA}}$, and each call appends only a finite number of possible terms to $\boldsymbol{t}$, the set of possible return values of $\mathcal{S}_{\mathrm{LRA}}$ is finite, and thus it is finite for $\varphi[\mathbf{k}, \mathbf{e}]$. ■

Lemma 7 *If $\mathcal{I}$ is a model for* LRA *and quantifier-free formula ψ, then $\mathcal{I}$ is also a model for $\psi\{\mathsf{e} \mapsto S_{R0}(\mathcal{I}, \psi, \mathsf{e})\}$.*

Proof: Let M be a set of literals of the form described in the definition of S_{R0} for $\mathcal{I}$, ψ and e. Consider the case where $S_{R0}(\mathcal{I}, \psi, \mathsf{e}) = \ell_i + \delta_i^\ell$ for some i, where $n > 0$. We show that $\mathcal{I}$ satisfies $M\{\mathsf{e} \mapsto \ell_i + \delta_i^\ell\}$. First, since $\mathsf{max}\{(\ell_1 + \delta_1^\ell)^\mathcal{I}, \ldots, (\ell_n + \delta_n^\ell)^\mathcal{I}\} = (\ell_i + \delta_i^\ell)^\mathcal{I}$, we know that $\mathcal{I}$ satisfies $M_\ell\{\mathsf{e} \mapsto \ell_i + \delta_i^\ell\}$. In the case that the bound on e we consider is strict, that is, $\mathsf{e} > \ell_i \in M_\ell$, then δ_i^ℓ is δ, and $\ell_i^\mathcal{I} < u_j^\mathcal{I}$ for all $j \in \{1, \ldots, m\}$. Thus, $\mathcal{I}$ satisfies $(\ell_i + \delta \prec u_j) = (\mathsf{e} \prec u_j)\{\mathsf{e} \mapsto \ell_i + \delta\}$. In the case that the bound on e we consider is non-strict, that is, if $\mathsf{e} \geq \ell_i \in M_\ell$, then δ_i^ℓ is 0, and $\ell_i^\mathcal{I} \leq u_j^\mathcal{I}$ for all $j \in \{1, \ldots, m\}$. Thus, $\mathcal{I}$ satisfies $(\ell_i \prec u_j) = (\mathsf{e} \prec u_j)\{\mathsf{e} \mapsto \ell_i\}$. In either case, we have that $\mathcal{I}$ satisfies each literal in $M_u\{\mathsf{e} \mapsto \ell_i + \delta_i^\ell\}$. Finally, $\mathcal{I}$ clearly satisfies $M_c\{\mathsf{e} \mapsto \ell_i + \delta_i^\ell\} = M_c$. The case when $m > 0$ is symmetric to the case when $n > 0$. In the case where $n = 0$ and $m = 0$, we have that ψ does not contain e, and

$\mathcal{I}$ satisfies $M\{\mathsf{e} \mapsto 0\}$. In each case, $\mathcal{I}$ satisfies $M\{\mathsf{e} \mapsto S_{R0}(\mathcal{I}, \psi, \mathsf{e})\}$, which entails $\psi\{\mathsf{e} \mapsto S_{R0}(\mathcal{I}, \psi, \mathsf{e})\}$, and thus the lemma holds. ■

Lemma 8 *$\mathcal{S}_{\mathrm{LRA}}$ is model-preserving for $\varphi[\mathbf{k}, \mathbf{e}]$.*

Proof: By the definition of S_R and repeated applications of Lemma 7. ■

Theorem 2 *$\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ is a sound and complete procedure for determining the* LRA*-satisfiability of $\exists \boldsymbol{a} \forall \boldsymbol{x}\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$.*

Proof: By Theorem 1 and Lemma 3 of our framework as well as LRA-specific Lemma 6 and Lemma 8. ■

We illustrate the procedure through examples. S_{R0} is non-deterministic; we choose instantiations only based on the lower bounds M_ℓ found in the procedure S_{R0}, though the procedure is free to base its instantiations on the upper bounds M_u as well. We underline the literal in M_ℓ corresponding to the bound whose value is maximal in $\mathcal{I}$. Γ is initially empty and on each iteration Γ' is the union of Γ and the skolemized negation of the input formula. Each round of $\mathcal{S}_{\mathrm{LRA}}$ computes a tuple $\boldsymbol{t}[\mathbf{k}]$, which is used to instantiate our quantified formula in Figure 1. The last column shows the corresponding instance of the quantified formula after simplification, including the elimination of δ.

Example 3. To demonstrate how non-strict bounds are handled, consider the formula $\forall x\,(x \leq a \vee x \leq b)$, whose skolemized negation is $e > a \wedge e > b$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ on this input is as follows.

			$S_{R0}(\mathcal{I}, \Gamma, \mathsf{e})$				
#	Γ	Γ'	e	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e	$\{\underline{e > a}, e > b\}$	$a + \delta$	$(a + \delta)$	$a < b$
2	sat	sat	e	$\{e > a, \underline{e > b}\}$	$b + \delta$	$(b + \delta)$	$b < a$
3	unsat						

This run shows $\exists ab \forall x\,(a \geq x \vee x \geq b)$ is LRA-unsatisfiable. The disjuncts of the instance $a + \delta \leq a \vee a + \delta \leq b$ added to Γ on the first iteration simplify to $\bot$ and $a < b$ respectively. We similarly obtain $b < a$ on the second iteration. ■

Example 4. To demonstrate how multiple universally quantified variables are handled, consider the formula $\forall xy\,(x + y < a \vee x - y < b)$ whose skolemized negation is $e_1 + e_2 \geq a \wedge e_1 - e_2 \geq b$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ is as follows.

			$S_{R0}(\mathcal{I}, \Gamma, \mathsf{e})$				
#	Γ	Γ'	e	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e_1	$\{e_1 \geq a - e_2, \underline{e_1 \geq b + e_2}\}$	$b + e_2$		
			e_2	$\{\underline{e_2 \geq \frac{a-b}{2}}\}$	$\frac{a-b}{2}$	$(\frac{a+b}{2}, \frac{a-b}{2})$	$\bot$
2	unsat						

This run shows $\exists ab \forall xy\,(x + y < a \vee x - y < b)$ is LRA-unsatisfiable. The substitution for e_2 is chosen based on M_ℓ after applying the substitution $\{e_1 \mapsto b + e_2\}$. ■

Example 5. To demonstrate how quantified formulas with Boolean structure are handled, consider the formula $\forall x\,((a < x \wedge x < b) \vee x < a + b)$ whose skolemized negation is $(a \geq e \vee e \geq b) \wedge e \geq a + b$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ is as follows.

			$S_{R0}(\mathcal{I}, \Gamma, \mathsf{e})$				
#	Γ	Γ'	e	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e	$\{\underline{e \geq a + b}\}$	$a + b$	$(a + b)$	$0 < b \wedge a < 0$
2	sat	sat	e	$\{\underline{e \geq b}, e \geq a + b\}$	b	(b)	$0 < a$
3	unsat						

This run shows $\exists ab\,\forall x\,((a < x \wedge x < b) \vee x < a + b)$ is LRA-unsatisfiable. On the first iteration, we assume that the propostionally satisfying assignment for Γ' included $a \geq e$, and hence $e \geq b$ is not included as a lower bound on that iteration. On the second iteration, the solver must satisfy both $b > 0$ and $a < 0$, which implies the model $\mathcal{I}$ is such that $(a + b)^{\mathcal{I}} > a^{\mathcal{I}}$ hence $e \geq b$ must exist in M_ℓ, and moreover $b^{\mathcal{I}} > (a + b)^{\mathcal{I}}$ hence b must be the maximal lower bound for e. ■

Example 6. To demonstrate a case where a variable has no bounds, consider the formula $\forall xy\,x \leq y$, whose skolemized negation is $e_1 > e_2$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ on this input is as follows.

			$S_{R0}(\mathcal{I}, \Gamma, \mathsf{e})$				
#	Γ	Γ'	e	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e_1	$\{\underline{e_1 > e_2}\}$	$e_2 + \delta$		
			e_2	$\emptyset$	0	$(\delta, 0)$	$\bot$
2	unsat						

This run shows $\forall xy\,x > y$ is LRA-unsatisfiable. Notice that after the substitution $\{e_1 \mapsto e_2 + \delta\}$, we have that Γ' contains neither an upper nor a lower bound for e_2, and hence we choose to return the value 0. ■

Example 7. To demonstrate a non-trivial case using the infinitesimal δ, consider the formula $\forall xy\,(x \leq 0 \vee y - 2 \cdot x \leq 0)$ whose skolemized negation is $e_1 > 0 \wedge e_2 - 2 \cdot e_1 > 0$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$ on this input is as follows.

			$S_{R0}(\mathcal{I}, \Gamma, \mathsf{e})$				
#	Γ	Γ'	e	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e_1	$\{\underline{e_1 > 0}\}$	δ		
			e_2	$\{\underline{e_2 > 2 \cdot \delta}\}$	$3 \cdot \delta$	$(\delta, 3 \cdot \delta)$	$\bot$
2	unsat						

This run shows $\forall xy\,x \leq 0 \vee y - 2 \cdot x \leq 0$ is LRA-unsatisfiable. ■

The procedure $\mathcal{P}_{\mathcal{S}_{\mathrm{LRA}}}$, which is an instance of the procedure in Figure 1, can be understood as lazily enumerating the disjuncts of the Loos-Weispfenning method for quantifier elimination over linear real arithmetic [33], with minor differences[4]. In this

[4] For instance, that method uses a distinguished term ∞ representing an arbitrarily large positive value.

$$\text{Return } \begin{cases} \frac{u_j+\ell_i}{2} & n>0 \text{ and } m>0 \\ \ell_i+1 & n>0 \text{ and } m=0 \\ u_j-1 & n=0 \text{ and } m>0 \\ 0 & n=0 \text{ and } m=0 \end{cases}, \text{ where } \begin{cases} \mathsf{max}\{\ell_1^{\mathcal{I}},\ldots,\ell_n^{\mathcal{I}}\}=\ell_i^{\mathcal{I}} \text{ if } n>0 \\ \mathsf{min}\{u_1^{\mathcal{I}},\ldots,u_m^{\mathcal{I}}\}=u_j^{\mathcal{I}} \text{ if } m>0. \end{cases}.$$

Fig. 4. An alternative return value for S_{R0}.

$$\text{Return one of } \begin{cases} l_j+\delta & n>0 \\ u_j-\delta & m>0 \\ \infty & m=0 \\ -\infty & n=0 \end{cases}, \text{ where } \begin{cases} \mathsf{max}\{\ell_1^{\mathcal{I}},\ldots,\ell_n^{\mathcal{I}}\}=\ell_i^{\mathcal{I}} \text{ if } n>0 \\ \mathsf{min}\{u_1^{\mathcal{I}},\ldots,u_m^{\mathcal{I}}\}=u_j^{\mathcal{I}} \text{ if } m>0. \end{cases}.$$

Fig. 5. An alternative return value for S_{R0} that is analogous to Loos and Weispfenning's method.

way, our approach is similar to the projection-based procedures described in [12, 29]. These approaches compute implicants of quantified formulas, while our approach instead computes a term which is in turn used for instantiation. This choice is important for our purposes, as it enables a uniform combination of the approach with existing instantiation-based techniques for first-order logic [17,22,49], and allows the approach to be used as a subprocedure for synthesis as described in [47].

An alternative return value for S_{R0} is presented in Figure 4. When using this return value, the procedure $\mathcal{P}_{\mathcal{S}_{\text{LRA}}}$ enumerates the disjuncts of Ferrante and Rackoff's method for quantifier elimination over linear real arithmetic [21], with minor differences. We provide an experimental evaluation of both of these selection functions in Section 6.

3.1 Comparison to Existing Approaches

As mentioned, at their core, most approaches for solving quantified linear arithmetic (including ours) share many similarities with one another. In particular, given an existentially quantified formula $\exists \boldsymbol{x}.\varphi$, based on some strategy, they enumerate (possibly lazily) a finite set of ground formulas that are entailed by this formula. We give a brief overview contrasting the technical details of existing approaches in this section.

We have mentioned that the approach in Figure 3 involves the use of a free distinguished constant δ, representating an infinitessmal positive value. Other approaches also involve use of a free distinguished constant ∞, representing an arbitrarily large positive value. Like δ, this term can be eliminated, as follows:

$$\infty < t \rightsquigarrow \bot \text{ and } \infty > t \rightsquigarrow \top \text{ where } \infty \notin FV(t).$$

The two most widely known algorithms for quantifier elimination for linear real arithmetic are the method based on infinitessimals in [33], and the method based on an

$$\text{Return } \begin{cases} \frac{u_j+\ell_i}{2} & n>0 \text{ and } m>0 \\ \infty & m=0 \\ -\infty & n=0 \end{cases} \text{, where } \begin{cases} \mathsf{max}\{\ell_1^{\mathcal{I}},\ldots,\ell_n^{\mathcal{I}}\}=\ell_i^{\mathcal{I}} \text{ if } n>0 \\ \mathsf{min}\{u_1^{\mathcal{I}},\ldots,u_m^{\mathcal{I}}\}=u_j^{\mathcal{I}} \text{ if } m>0. \end{cases}$$

Fig. 6. An alternative return value for S_{R0} that is analogous to Ferrante and Rackoff's method.

interior point method in [21]. To put these algorithms into the context of our approach, we provide two additional alternatives for the return value of S_{R0} (Figures 5 and 6) that closely approximate the effect of these methods.

Recent approaches are inspired by of one (or both) of these methods. The approaches described in [12, 29] are closely based on the Loos-Weispfenning method, and the approach described in [18] is closely based on Ferrante-Rackoff method. The approach described in [40] examines a certified version of both approaches.

As mentioned, Figure 3 is inspired by the Loos-Weispfenning method, but does not use infinities. Similarly, the return value in Figure 4 is inspired by Ferrante-Rackoff method, but does not use infinities. One possible advantage of the approach in Figure 4 is in the context of quantifier alternation. In particular, that selection function does not use any virtual terms, and thus we may consider instances of quantified formulas having nested quantification where eliminating virtual terms is not obvious. A complete strategy for quantifier alternation using this selection function is the subject of future work.

4 Instantiation for Quantifier-Free LIA-Formulas

We now turn our attention to the class of arbitrary LIA-formulas $\exists \boldsymbol{a}\, \forall x\, \varphi[\boldsymbol{a}, \boldsymbol{x}]$, $\boldsymbol{x}$ and $\boldsymbol{a}$ are vectors of Int variables, and where $\varphi[\boldsymbol{a}, \boldsymbol{x}]$ is quantifier-free. We again assume all equalities are eliminated from φ by replacing them with a conjunction of inequalities.

Figure 7 gives a selection function $\mathcal{S}_{\mathrm{LIA}}$ for LIA. The procedure invokes the recursive procedure S_I, which takes as arguments $\mathcal{I}$, $\neg\varphi[\mathbf{k}, \mathbf{e}]$, variables $\mathbf{e}$ that we have yet to incorporate into the substitutions, an integer θ, terms $\boldsymbol{t}$ found as substitutions for variables from $\mathbf{e}$ so far, and a tuple of symbols $\boldsymbol{p}$ from $\{+, -\}$ which we refer to as *polarities*. The role of θ will be to capture divisibility relationships through the procedure, where θ is initially 1. The procedure invokes a call to $S_{I0}(\mathcal{I}, \psi, \mathsf{e}_i)$ which based on the propositionally satisfying assignment for ψ returns a tuple of the form (c, t_i, p_i), where c is a constant, t_i is a term, and p_i is a polarity. The procedure for constructing the term t_i in the procedure S_{I0} is similar to the procedure S_{R0} in the previous section, where we find the lower bound of the form $c_i \cdot \mathsf{e} \geq \ell_i$ such that the (rational) value $(\frac{\ell_i}{c_i})^{\mathcal{I}}$ is maximal, and similarly for M_u. Additionally, S_{R0} adds a constant ρ to the maximal lower bound (resp. minimal lower bound). This constant ensures that the returned term t_i and e are congruent modulo $\theta \cdot c$ in $\mathcal{I}$, a fact which in part suffices to show the overall function to be model-preserving. It then constructs a *substitution with coefficients* σ of the form $\{c \cdot \mathsf{e}_i \mapsto t_i\}$. A substitution of this form may be applied to integer terms of the

$\mathcal{S}_{\rm LIA}(\mathcal{I}, \Gamma, \neg\varphi[\mathsf{k}, \mathsf{e}])$:

Return $S_I(\mathcal{I}, \neg\varphi[\mathsf{k}, \mathsf{e}], \mathsf{e}, 1, (), ())$

$S_I(\mathcal{I}, \psi, (\mathsf{e}_i, \ldots, \mathsf{e}_n), \theta, \boldsymbol{t}, \boldsymbol{p})$:

If $i > n$, return $\boldsymbol{t}\ \mathsf{div}^{\boldsymbol{p}}\ \theta$
Otherwise, let $(c, t_i, p_i) = S_{I0}(\mathcal{I}, \psi, \mathsf{e}_i, \theta)$, $\sigma = \{c \cdot \mathsf{e}_i \mapsto t_i\}$
Return $S_I(\mathcal{I}, \psi\sigma, (\mathsf{e}_{i+1}, \ldots, \mathsf{e}_n), \theta \cdot c, ((c \cdot \boldsymbol{t})\sigma, \theta \cdot t_i), (\boldsymbol{p}, p_i))$

$S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta)$:

Let $M = M_\ell \cup M_u \cup M_c$ be such that:
- $\mathcal{I} \models M$ and $M \models_p \psi$,
- $M_\ell \Leftrightarrow \{c_1 \cdot \mathsf{e} \geq \ell_1, \ldots, c_n \cdot \mathsf{e} \geq \ell_n\}$, $c_1 > 0, \ldots, c_n > 0$,
- $M_u \Leftrightarrow \{d_1 \cdot \mathsf{e} \leq u_1, \ldots, d_m \cdot \mathsf{e} \leq u_m\}$, $d_1 > 0, \ldots, d_m > 0$, and
- $\mathsf{e} \notin FV(\ell_1, \ldots, \ell_n) \cup FV(u_1, \ldots, u_m) \cup FV(M_c)$.

Return one of
$$\begin{cases} (c_i, \ell_i + \rho, +) & n > 0, \mathsf{max}\{(\frac{\ell_1}{c_1})^{\mathcal{I}}, \ldots, (\frac{\ell_n}{c_n})^{\mathcal{I}}\} = (\frac{\ell_i}{c_i})^{\mathcal{I}}, \\ & \rho = (c_i \cdot \mathsf{e} - \ell_i)^{\mathcal{I}}\ \mathsf{mod}\ (\theta \cdot c_i) \\ (d_j, u_j - \rho, -) & m > 0, \mathsf{min}\{(\frac{u_1}{d_1})^{\mathcal{I}}, \ldots, (\frac{u_m}{d_m})^{\mathcal{I}}\} = (\frac{u_j}{d_j})^{\mathcal{I}}, \\ & \rho = (u_j - d_j \cdot \mathsf{e})^{\mathcal{I}}\ \mathsf{mod}\ (\theta \cdot d_j) \\ (1, \rho, +) & n = 0, m = 0, \rho = \mathsf{e}^{\mathcal{I}}\ \mathsf{mod}\ \theta \end{cases}$$

Fig. 7. A selection function $\mathcal{S}_{\rm LIA}$ for arbitrary quantifier-free LIA-formula $\varphi[\boldsymbol{a}, x]$.

form $c \cdot (d \cdot \mathsf{e}_i + s)$ where $\mathsf{e}_i \notin FV(s)$, where $(c \cdot (d \cdot \mathsf{e}_i + s))\sigma$ is defined as $d \cdot t_i + c \cdot s$. Additionally, we define $(s_1 \bowtie s_2)\sigma$ as $(c \cdot s_1)\sigma \bowtie (c \cdot s_2)\sigma$ for $\bowtie \in \{<, >\}$, and thus we can apply σ to arbitrary LIA-formulas. After constructing σ, the procedure S_I invokes a recursive call where σ is applied to ψ and $(c \cdot \boldsymbol{t})$, θ is multiplied by c, the term $\theta \cdot t_i$ is appended to $\boldsymbol{t}$, and p_i is appended to $\boldsymbol{p}$.

Overall, S_I returns a vector of terms $(\boldsymbol{t}\ \mathsf{div}^{\boldsymbol{p}}\ \theta)$, that is, integer division applied pairwise to the terms in $\boldsymbol{t}$ and the constant θ, where $\boldsymbol{p}$ determines whether this division rounds up or down. We add the instance $\varphi[\mathbf{k}, \boldsymbol{t}\ \mathsf{div}^{\boldsymbol{p}}\ \theta] \downarrow$ to Γ in Figure 1, where occurrences of integer division are eliminated by defining $\varphi[\mathbf{k}, \boldsymbol{t}\ \mathsf{div}^{\boldsymbol{p}}\ \theta] \downarrow$ as $\varphi[\mathbf{k}, \mathbf{d}] \wedge \theta \cdot \mathbf{d} \approx \boldsymbol{t} \pm^{\boldsymbol{p}} \mathbf{m} \wedge 0 \leq \mathbf{m} < \theta$, where $\mathbf{d}$ and $\mathbf{m}$ are distinct fresh constants, and $\pm^p$ is $+$ if p is $+$ and analogously for $-$. Note that our selection function chooses $\boldsymbol{p}$ such that integer division rounds up for terms coming from lower bounds, and rounds down for terms combing from upper bounds. This choice is not required for correctness, but can reduce the number of instances needed for showing unsatisfiability.

Lemma 9 *$\mathcal{S}_{\rm LIA}$ is finite for $\varphi[\mathsf{k}, \mathsf{e}]$.*

Proof: First, we show that only a finite number of tuples are returned by $S_{I0}(\mathcal{I}, (\neg\varphi[\mathsf{k}, \mathsf{e}])\sigma, \mathsf{e}_i, \theta)$ for any $\mathcal{I}, \sigma, \mathsf{e}_i$ and finite θ. Let A be the set of atoms occurring in $\varphi[\mathsf{k}, \mathsf{e}_i]\sigma$. The literals in satisfying assignments of $(\neg\varphi[\mathsf{k}, \mathsf{e}])\sigma$ are over these

atoms. Let $\{c_1 \cdot \mathsf{e}_i < \ell_1, \ldots, c_n \cdot \mathsf{e}_i < \ell_n, d_1 \cdot \mathsf{e}_i > u_1, \ldots, d_m \cdot \mathsf{e}_i > u_m\}$ be the set of atoms that are in solved form with respect to e_i that are equivalent to the atoms of A containing e_i, where $\mathsf{e}_i \notin FV(\ell_1, \ldots, \ell_n, u_1, \ldots, u_m)$ and $c_1 > 0, \ldots, c_n > 0, d_1 > 0, \ldots, d_m > 0$. The tuples returned by $S_{I0}(\mathcal{I}, (\neg\varphi[\mathbf{k}, \mathbf{e}])\sigma, \mathsf{e}_i)$ are in the finite set:

$$\begin{array}{llllll} \{(c_i, \ell_i + \rho, +) \mid 1 \leq i \leq n\} & \cup & \{(d_j, u_j + 1 + \rho, -) \mid 1 \leq j \leq m\} & \cup & \\ \{(d_j, u_j - \rho, -) \mid 1 \leq j \leq m\} & \cup & \{(c_i, \ell_i - 1 - \rho, +) \mid 1 \leq i \leq n\} & \cup & \{(1, \rho, +)\} \end{array}$$

and where $0 \leq \rho < (\theta \cdot c)$. Since c and θ are finite, there are a finite number of tuples of this form. Since there are only a finite number of recursive calls to S_I within $\mathcal{S}_{\mathrm{LIA}}$, and each call modifies $\boldsymbol{t}$ based a finite number of possible tuples coming from the set above, the set of possible return values of $\mathcal{S}_{\mathrm{LIA}}$ is finite, and thus it is finite for $\varphi[\mathbf{k}, \mathbf{e}]$. ■

Lemma 10 *If $\mathcal{I}$ is a model for* LIA *and a quantifier-free formula ψ, $\theta \geq 1$, and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (c, t, p)$, then:*

1. $(c \cdot \mathsf{e})^{\mathcal{I}} \equiv_{\theta \cdot c} t^{\mathcal{I}}$, *and*
2. $\mathcal{I} \models \psi\{c \cdot \mathsf{e} \mapsto t\}$.

Proof: We first show part 1. In the case that $n > 0$ and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (c_i, \ell_i + \rho, +)$, and

$$(\ell_i + \rho)^{\mathcal{I}} \equiv_{\theta \cdot c_i} (\ell_i + (c_i \cdot \mathsf{e} - \ell_i)^{\mathcal{I}} \bmod (\theta \cdot c_i))^{\mathcal{I}} \equiv_{\theta \cdot c_i} (c_i \cdot \mathsf{e})^{\mathcal{I}}.$$

In the case that $m > 0$ and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (d_j, u_j - \rho, -)$, we have

$$(u_j - \rho)^{\mathcal{I}} \equiv_{\theta \cdot d_j} (u_j - (u_j - d_j \cdot \mathsf{e})^{\mathcal{I}} \bmod (\theta \cdot d_j))^{\mathcal{I}} \equiv_{\theta \cdot d_j} (d_j \cdot \mathsf{e})^{\mathcal{I}}.$$

In the case that $n = 0$, $m = 0$, and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (1, \rho, +)$, we have that $\rho^{\mathcal{I}} \equiv_{\theta \cdot 1} (\mathsf{e}^{\mathcal{I}} \bmod \theta)^{\mathcal{I}} \equiv_{\theta \cdot 1} (1 \cdot \mathsf{e})^{\mathcal{I}}$.

To show part 2, we first focus on the case where $n > 0$ and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (c_i, \ell_i + \rho, +)$. We have that $\rho = (c_i \cdot \mathsf{e} - \ell_i)^{\mathcal{I}} \bmod (\theta \cdot c_i)$. Let M be a set of literals of the form described in the body of $S_{I0}(\mathcal{I}, \psi, \mathsf{e})$. We show that $\mathcal{I}$ satisfies each literal in $M\sigma$, where $\sigma = \{c_i \cdot \mathsf{e} \mapsto \ell_i + \rho\}$. First, consider an atom in $M_\ell\sigma$ that is equivalent to $(c_j \cdot \mathsf{e} \geq \ell_j)\sigma$ for some $j \in \{1, \ldots, n\}$. This is equivalent to $(c_j \cdot c_i \cdot \mathsf{e} \geq c_i \cdot \ell_j)\sigma$, which is equivalent to $\frac{c_j \cdot c_i}{c_i} \cdot (\ell_i + \rho) \geq \frac{c_j \cdot c_i}{c_j} \cdot \ell_j$, which is satisfied by $\mathcal{I}$ since $(\frac{\ell_i}{c_i})^{\mathcal{I}} \geq (\frac{\ell_j}{c_j})^{\mathcal{I}}$ by our selection of $(c_i, \ell_i + \rho)$ and since $\rho \geq 0$. Second, consider the atom in $M_u\sigma$ that is equivalent to $(d_j \cdot \mathsf{e} \leq u_j)\sigma$ for some $j \in \{1, \ldots, m\}$. Let $\rho' = (c_i \cdot \mathsf{e} - \ell_i)^{\mathcal{I}}$, which is greater than 0 since $\mathcal{I}$ satisfies $(c_i \cdot \mathsf{e} \geq \ell_i)$. Since $(c_i \cdot \mathsf{e})^{\mathcal{I}} = (\ell_i + \rho')^{\mathcal{I}}$, we have that $\mathcal{I}$ satisfies $(d_j \cdot \mathsf{e} \leq u_j)\{c_i \cdot \mathsf{e} \mapsto \ell_i + \rho'\}$, which is equivalent to $(d_j \cdot (\ell_i + \rho') \leq c_i \cdot u_j)$. Since $\rho = \rho' \bmod (\theta \cdot c) \leq \rho'$, we have that $\mathcal{I}$ also satisfies $(d_j \cdot (\ell_i + \rho) \leq c_i \cdot u_j)$, which is $(d_j \cdot \mathsf{e} \leq u_j)\sigma$. Finally, $\mathcal{I}$ satisfies $M_c\sigma$ as $M_c\sigma = M_c$ and $\mathcal{I} \models M_c$. Thus, $\mathcal{I}$ satisfies $M\sigma$, which entails $\psi\sigma$. The case for when $m > 0$ and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}, \theta) = (d_j, u_j - \rho, -)$ is symmetric. When $n = 0$, $m = 0$, and $S_{I0}(\mathcal{I}, \psi, \mathsf{e}) = (1, \rho, +)$, the assignment M does not contain e, and thus $\mathcal{I}$ satisfies $M\{c \cdot \mathsf{e} \mapsto \rho\} = M$ and $\psi\{c \cdot \mathsf{e} \mapsto \rho\}$. ■

Lemma 11 *Each recursive call to $S_I(\mathcal{I}, \psi, (\mathsf{e}_i, \ldots, \mathsf{e}_n), \theta, (t_1, \ldots, t_{i-1}), \boldsymbol{p})$ within $\mathcal{S}_{\mathrm{LIA}}(\mathcal{I}, \Gamma, (\mathsf{e}_1, \ldots, \mathsf{e}_n))$ is such that:*

1. $\theta \mid t_j^{\mathcal{I}}$ *for each* $1 \leq j < i$, *and*
2. $\mathcal{I} \models \psi$ *and* ψ *is equivalent to* $\neg\varphi[\mathbf{k}, \mathbf{e}]\{\theta \cdot \mathsf{e}_1 \mapsto t_1\} \cdot \ldots \cdot \{\theta \cdot \mathsf{e}_{i-1} \mapsto t_{i-1}\}$.

Proof: Both statements clearly hold for the initial call to S_I in the body of $\mathcal{S}_{\mathrm{LIA}}$. Now, assume both statements hold for some call to $S_I(\mathcal{I}, \psi, (\mathsf{e}_i, \mathbf{e}'), \theta, (t_1, \ldots, t_{i-1}), \boldsymbol{p})$, and assume $(c, t_i, p_i) = S_{I0}(\mathcal{I}, \psi, \mathsf{e}_i, \theta)$. We show that both statements hold for the call to $S_I(\mathcal{I}, \psi\sigma, \mathbf{e}', \theta \cdot c, ((c \cdot t_1)\sigma, \ldots, (c \cdot t_{i-1})\sigma, \theta \cdot t_i), (\boldsymbol{p}, p_i))$, where $\sigma = \{c \cdot \mathsf{e}_i \mapsto t_i\}$.

To show part 1, we have from Lemma 10 part 1 that:

$$(c \cdot \mathsf{e}_i)^{\mathcal{I}} \equiv_{\theta \cdot c} (t_i + \rho)^{\mathcal{I}} \tag{7}$$

Consider a t_j where $1 \leq j < i$, where by our assumption is such that $\theta \mid t_j^{\mathcal{I}}$, and thus $\theta \cdot c \mid (c \cdot t_j)^{\mathcal{I}}$. By (7), we have that $\theta \cdot c \mid ((c \cdot t_j)\sigma)^{\mathcal{I}}$. Also by (7), we have that $c \mid (t_i + \rho)^{\mathcal{I}}$, and thus $\theta \cdot c \mid (\theta \cdot (t_i + \rho))^{\mathcal{I}}$.

To show part 2, by our assumption, $\mathcal{I} \models \psi$ and thus by Lemma 10 part 2 we have that $\mathcal{I} \models \psi\sigma$. By our assumption, ψ is equivalent to $\neg\varphi[\mathbf{k}, \mathbf{e}]\{\theta \cdot \mathsf{e}_1 \mapsto t_1\} \cdot \ldots \cdot \{\theta \cdot \mathsf{e}_{i-1} \mapsto t_{i-1}\}$. Thus, $\psi\sigma$ is equivalent to $\neg\varphi[\mathbf{k}, \mathbf{e}]\{(\theta \cdot c) \cdot \mathsf{e}_1 \mapsto (c \cdot t_1)\sigma\} \cdot \ldots \cdot \{(\theta \cdot c) \cdot \mathsf{e}_{i-1} \mapsto (c \cdot t_{i-1})\sigma\} \cdot \{(\theta \cdot c) \cdot \mathsf{e}_i \mapsto \theta \cdot (t_i + \rho)\}$. Thus, the lemma holds. ■

Lemma 12 $\mathcal{S}_{\mathrm{LIA}}$ *is model-preserving for* $\varphi[\mathbf{k}, \mathbf{e}]$.

Proof: Assume that $\mathcal{S}_{\mathrm{LIA}}(\mathcal{I}, \Gamma, \neg\varphi[\mathbf{k}, \mathbf{e}]) = \boldsymbol{t}$, where $\mathbf{e} = (\mathsf{e}_1, \ldots, \mathsf{e}_n)$, and $\boldsymbol{t} = (t_1, \ldots, t_n)$. By Lemma 11 and the definition of $\mathcal{S}_{\mathrm{LIA}}$, there is a θ such that for each $i = 1, \ldots, n$, term t_i is of the form $s_i \ \mathsf{div}^p \ \theta$ where $\theta \mid s_i^{\mathcal{I}}$, and $\mathcal{I} \models (\neg\varphi[\mathbf{k}, \mathbf{e}])\{\theta \cdot \mathsf{e}_1 \mapsto s_1\} \cdot \ldots \cdot \{\theta \cdot \mathsf{e}_n \mapsto s_n\}$. Thus, $\mathcal{I}$ satisfies $(\neg\varphi[\mathbf{k}, \mathbf{e}])\{\mathbf{e} \mapsto \boldsymbol{t}\} = \neg\varphi[\mathbf{k}, \boldsymbol{t}]$, and thus $\mathcal{S}_{\mathrm{LIA}}$ is model-preserving for $\varphi[\mathbf{k}, \mathbf{e}]$. ■

Theorem 3 $\mathcal{P}_{\mathcal{S}_{\mathrm{LIA}}}$ *is a sound and complete procedure for determining the* LIA*-satisfiability of* $\exists \boldsymbol{a} \, \forall \boldsymbol{x} \, \varphi[\boldsymbol{a}, \boldsymbol{x}]$.

Proof: By Theorem 1, Lemma 3, Lemma 9 and Lemma 12. ■

Example 8. To demonstrate a case involving a substitution with coefficients, consider the formula $\forall xy \, (2 \cdot x < a \vee x + 3 \cdot y < b)$ whose negation is $2 \cdot e_1 \geq a \wedge e_1 + 3 \cdot e_2 \geq b$. A possible run of $\mathcal{P}_{\mathcal{S}_{\mathrm{LIA}}}$ on this input is as follows.

#	Γ	Γ'	$S_{I0}(\mathcal{I}, \Gamma, \mathsf{e}, \theta)$ e	θ	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e_1	1	$\{\underline{2 \cdot e_1 \geq a}, \ldots\}$	$(2, a, +)$		
			e_2	2	$\{\underline{6 \cdot e_2 \geq 2 \cdot b - a}\}$	$(6, 2 \cdot b - a, +)$	$(6 \cdot a, 4 \cdot b - 2 \cdot a) \ \mathsf{div}^+ \ 12$	ψ_1
2	unsat							

Thus, $\exists ab \, \forall xy \, (2 \cdot x < a \vee x + 3 \cdot y < b)$ is LIA-unsatisfiable. We assume $\rho = 0$ for all calls to S_{I0} in this run. Applying the substitution $\{2 \cdot e_1 \mapsto a\}$ to $e_1 + 3 \cdot e_2 \geq b$ results in the bound $6 \cdot e_2 \geq 2 \cdot b - a$ for e_2. We add to Γ the instance ψ_1, which is equivalent to $2 \cdot ((6 \cdot a) \ \mathsf{div}^+ \ 12) < a \vee (6 \cdot a) \ \mathsf{div}^+ \ 12 + 3 \cdot ((4 \cdot b - 2 \cdot a) \ \mathsf{div}^+ \ 12) < b$. Applying normalization $\downarrow$ to this formula results in a one that is LIA-unsatisfiable. ■

Example 9. To demonstrate a case involving a non-zero value of ρ, consider the formula $\forall xy\,(3\cdot x+y \not\approx a \vee 0 > y \vee y > 2)$ whose negation is $3\cdot e_1+e_2 \approx a \wedge 0 \leq e_2 \wedge e_2 \leq 2$, where $\approx$ denotes the conjunction of non-strict upper and lower bounds. A possible run of $\mathcal{P}_{\mathcal{S}_{\text{LIA}}}$ on this input is as follows.

			$S_{I0}(\mathcal{I}, \Gamma, \mathsf{e}, \theta)$					
#	Γ	Γ'	e	θ	M_ℓ	return	$\boldsymbol{t}[\mathbf{k}]$	Add to Γ
1	sat	sat	e_1	1	$\{\underline{3\cdot e_1 \geq a - e_2}\}$	$(3, a-e_2, +)$		
			e_2	3	$\{\underline{e_2 \geq 0}\}$	$(1, 0, +)$	$(a, 0)\ \mathsf{div}^+\ 3$	ψ_1
2	sat	sat	e_1	1	$\{\underline{3\cdot e_1 \geq a - e_2}\}$	$(3, a-e_2, +)$		
			e_2	3	$\{\underline{e_2 \geq 0}\}$	$(1, 1, +)$	$(a-1, 1)\ \mathsf{div}^+\ 3$	ψ_2
3	sat	sat	e_1	1	$\{\underline{3\cdot e_1 \geq a - e_2}\}$	$(3, a-e_2, +)$		
			e_2	3	$\{\underline{e_2 \geq 0}\}$	$(1, 2, +)$	$(a-2, 2)\ \mathsf{div}^+\ 3$	ψ_3
4	unsat							

This run shows $\exists a\,\forall xy\,(2\cdot x < a \vee x + 3\cdot y < b)$ is LIA-unsatisfiable. On the first iteration, we assume that Γ' is satisfied by a model, call it $\mathcal{I}_1$, that interprets all variables as 0, and hence the values chosen for e_1 and e_2 correspond to their maximal lower bounds in $\mathcal{I}_1$, $a - e_2$ and 0 respectively, where in each call to S_{I0} we have $\rho = 0$. The instance ψ_1 added to Γ on this iteration is equivalent to $3\cdot(a\ \mathsf{div}^+\ 3) \not\approx a$ and implies that $a^{\mathcal{I}} \not\equiv_3 0$ in subsequent models $\mathcal{I}$. Thus, models $\mathcal{I}$ satisfying $3\cdot e_1 + e_2 \approx a$ are such that $e_2^{\mathcal{I}} \not\equiv_3 0$. On the next iteration, Γ' is satisfied by a model, call it $\mathcal{I}_2$, where the maximal lower bound for e_2 is 0. By the above reasoning and since $\mathcal{I}_2$ satisfies $3\cdot e_1 + e_2 \approx a$, it must be that $\rho = ((e_2 - 0)^{\mathcal{I}_2}\ \mathsf{mod}\ 3) \neq 0$. Assume $(e_2 - 0)^{\mathcal{I}_2} \equiv_3 1$. The instance ψ_2 is equivalent to $3\cdot((a-1)\ \mathsf{div}^+\ 3) + 1 \not\approx a$, which implies that $a^{\mathcal{I}} \not\equiv_3 1$ in subsequent models $\mathcal{I}$, and hence $e_2^{\mathcal{I}} \not\equiv_3 1$. The instance ψ_3 is equivalent to $3\cdot((a-2)\ \mathsf{div}^+\ 3) + 2 \not\approx a$ and implies that $a^{\mathcal{I}} \not\equiv_3 2$, which together with the two previous instances are T-unsatisfiable. ■

The procedure $\mathcal{P}_{\mathcal{S}_{\text{LIA}}}$ can be understood to lazily enumerating disjuncts of Cooper's algorithm for quantifier elimination over linear integer arithmetic [15], with minor differences. The algorithm is essentially enumerating a single path of [15] by using the model to select a satisfied case split for each variable over an entire block of quantifiers.[5] Like that approach, the worst-case performance is dependent upon the size of coefficients of monomials, which is manifested in our case by the fact that the number of possible return values of S_{I0} is proportional to the size of θ. While not shown here, our implementation takes steps to reduce the size of θ by factoring out common divisors in θ and the coefficients returned by S_{I0}.

4.1 Comparison to Existing Approaches

The approach taken in $\mathcal{P}_{\mathcal{S}_{\text{LIA}}}$ is similar to the one taken in Section 2.5 in [12]. The most substantial difference between the two algorithms is that $\mathcal{P}_{\mathcal{S}_{\text{LIA}}}$ implements a variant of Cooper's algorithm while *resolve* in [12] uses the model to guide an execution of

[5] In the parlance of [15], $\mathcal{P}_{\mathcal{S}_{\text{LIA}}}$ selects a feasible j value using the calculation of ρ and avoids introducing the $F_{\pm\infty}$ cases by introducing the no bounds case ($n = 0, m = 0$) and always favoring bounds when one exists.

$\mathsf{SMTQI}(T, \Gamma)$:

If Γ is T-unsatisfiable, then return "unsat".
Otherwise,
  Let M be such that $\mathcal{I} \models M$ and $M \models_p \Gamma$ for some model $\mathcal{I}$ of T and Γ.
  $\Gamma' := \Gamma$.
  For each $(\neg)A_i \in M$ where $A_i \Leftrightarrow \forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$,
    $\Gamma' := \Gamma' \cup (A_i \vee B_i) \cup (B_i \Rightarrow \neg\varphi_i[\boldsymbol{e}_i])$.
    If $A_i \in M$ and $B_i \in M$,
      $\Gamma' := \Gamma' \cup (A_i \Rightarrow \varphi_i[\mathcal{S}_i(\mathcal{I}, \Gamma, \neg\varphi_i[\boldsymbol{e}_i])])$
  If $\Gamma' = \Gamma$, then return "sat". Otherwise, return $\mathsf{SMTQI}(T, \Gamma')$.

Fig. 8. Procedure SMTQI for SMT solving with quantifier instantiation, which determines the T-satisfiability of a ground set of T-formulas Γ in purified form with respect to quantified formulas.

the Omega test [44]. The most similar aspects of the approaches are the computation of a feasible ρ and the computation of the d values in the *grey* shadow cases of $resolve$. These differ in that a different d value is selected to ensure separation between each upper bound and the greatest lower bound in a *projection* whereas ρ is selected using the current value of $c_i \cdot e$ (the selection of d is agnostic to e in our parlance) to ensure all bounds are satisfied by a single instantiation.

5 Integration in an SMT Solver

This section gives an overview of how the instantiation-based procedure for quantified formulas as described in Section 2 can be integrated into a solving architecture used by SMT solvers, and used in part for determining the satisfiability of inputs Γ_0 having arbitrary Boolean structure that contain any number of quantified T-formulas.

Figure 8 defines a procedure SMTQI which takes as input a theory T and a set of ground T-formulas Γ in *purified form* with respect to quantified formulas, that is, a formula obtained from Γ_0 by replacing all quantified formulas $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$ in Γ_0 by (uniquely associated) boolean variables A_i. We call such a variable A_i the *positive guard* of $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$ and write $A_i \Leftrightarrow \forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$ to denote A_i is the positive guard of $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$. In each call to SMTQI, if Γ is T-unsatisfiable, then the procedure returns "unsat". Otherwise, we find a model $\mathcal{I}$ of T and Γ and a corresponding propositionally satisfying assignment M. We then build a new set of formulas Γ', initially containing Γ, as follows. For each quantified formula $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$ whose positive guard A_i is in M, we add the formulas $(A_i \vee B_i)$ and $(B_i \Rightarrow \neg\varphi_i[\boldsymbol{e}_i])$ to Γ' if we have not done so already, where B_i is a fresh Boolean variable, which we call the *negative guard* of $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$. If both the positive and negative guards of $\forall \boldsymbol{x}.\varphi_i[\boldsymbol{x}]$ are asserted positively in M (we will say such a quantified formula is *active in* M), we consider an instance of this quantified formula to Γ' based on its associated selection function $\mathcal{S}_i$, If φ_i has nested quantification, this instance will contain quantified formulas, which we purify in

the same manner described above. If no new formulas are added to Γ' in this process, then the procedure returns "sat". Otherwise, we call SMTQI on Γ'.

At a high level, the procedure in Figure 8 adds guarded instances of quantified formulas to an evolving set of formulas Γ until Γ is T-unsatisfiable, or a fixed point is reached. In this respect, the algorithm is similar to existing instantiation-based approaches used by SMT solvers for quantified formulas [16, 23]. However, the procedure differs from these approaches in the following ways. Firstly, when a universally quantified formula are asserted negatively in M, typical approaches add the clause $(\neg A_i \Rightarrow \neg\varphi_i[\boldsymbol{e}_i])$ to Γ. Here, we choose to add the clauses $(A_i \vee B_i)$ and $(B_i \Rightarrow \neg\varphi_i[\boldsymbol{e}_i])$ instead. This allow us to consider both the positive and negative versions of quantified formulas simultaneously. As such, the algorithm only adds instances of quantified formulas where *both* the positive and negative guards are asserted positively in M. To ensure the model soundness of the approach (that is, the algorithm answers "sat" only if the input is indeed T-satisfiable), we require the following property of the set of literals M in the body of SMTQI:

If no quantified formula is active in M, then $\Gamma \cup \{B_i\}$ is T-unsat for $i = 1, \ldots, n$, where $\{A_1, \ldots, A_n\}$ is the set of the positive guards that are asserted positively in M.

In other words, the set of literals M are chosen such that, if possible, at least one of $B_1, \ldots, B_n$ is true in M. In practice, this requirement can be met in a DPLL(T)-based SMT solver by instructing its underlying SAT solver, when it chooses a decision literal to add to M, to choose one of the unassigned negative guards of quantified formulas in Γ, if one exists, and assert it positively.

We refer to the treatment of quantified formulas in Figure 8 as *counterexample-guided quantifier instantiation* [47]. A closely related approach is that of model-based quantifier instantiation [24], which like the approach described here, adds instances of quantified formulas based on models for their negations. This approach differs in its scope, in that it primarily targets quantified formulas having uninterpreted functions, whereas the approach described in Figure 8 targets quantified formulas having no uninterpreted functions. It also differs in that it uses a separate copy of the SMT solver as an oracle for checking the satisfiability of the negation of each quantified formula it instantiates, whereas the approach described in Figure 8 uses a single instance of the SMT solver for doing these tasks simultaneously in its main solving loop.

5.1 Arbitrary Quantifier Alternation

The algorithm in Figure 8 can be used as a basis for handling quantified formulas with arbitrary quantifier alternations. Assume we rewrite formulas added to Γ' in the body SMTQI so they are in purified form with respect to quantified formulas, that is, we replace each quantified formula $\forall \boldsymbol{x}.\varphi_j[\boldsymbol{x}]$ occurring in formulas $B_i \Rightarrow \neg\varphi_i[\boldsymbol{e}_i]$ and $A_i \Rightarrow \varphi_i[\mathcal{S}_i(\mathcal{I}, \Gamma, \boldsymbol{e}_i)]$ with its corresponding positive guard A_j. We give an intuition of how the procedure SMTQI handles such formulas in the following. A more comprehensive description is the subject of future work.

Consider the LIA-formula $\forall x.\neg(\forall y.x > y)$, call it φ_1, whose positive guard is A_1. Given the input $\Gamma = \{A_1\}$, the procedure SMTQI adds the formulas $A_1 \vee B_1$ and

$B_1 \Rightarrow A_2$ to Γ, where A_2 is the positive guard for $(\forall y.e_1 > y)$ (call this formula φ_2) where e_1 is a fresh constant. On the second call to SMTQI, the satisfying assignment includes A_2 and the procedure similarly adds the formulas $A_2 \vee B_2$ and $B_2 \Rightarrow e_1 \leq e_2$ to Γ where e_2 is a fresh constant. On the third call to SMTQI, we have that $\Gamma = \{A_1, B_1 \Rightarrow A_2, B_2 \Rightarrow e_1 \leq e_2, \ldots\}$ is T-satisfiable, and we may choose a satisfying assignment $M = \{A_1, B_1, A_2, B_2, e_1 \leq e_2, \ldots\}$. Both φ_1 and φ_2 are active in M. The literal $e_1 \leq e_2$ is over the atoms of $\varphi_2[e_2/y]$, and we may add the formula $A_2 \Rightarrow e_1 > e_1$ to Γ on this iteration, assuming our selection function for φ_2 chose to return the maximal lower bound e_1 for e_2. On the fourth call to SMTQI, we have that $\Gamma = \{A_1, B_1 \Rightarrow A_2, A_2 \Rightarrow e_1 > e_1, \ldots\}$, and hence B_1 or A_2 cannot be asserted positively in a satisfying assignment M for this set. Hence, neither φ_1 nor φ_2 is active in M and the procedure SMTQI adds no instances to Γ, indicating that our input is satisfiable.

6 Experimental Evaluation

We have implemented the procedure in the SMT solver CVC4 [4] (version 1.5 prerelease). This section presents an evaluation of this implementation compared against other SMT solvers and first-order theorem provers.

Pure Quantified Linear Arithmetic We considered all quantified benchmarks over 6 classes in the LRA and LIA logics of the SMT library [5]. The class **keymaera** are verification conditions coming from the Keymaera verification tool [42], **scholl** were used for simplification of non-convex polyhedra in [51], **psyco** were used for weakest precondition synthesis for compiler optimizations in [34], **uauto** correspond to verification conditions in [26], and the **tptp** classes correspond to simple arithmetic conjectures coming from the TPTP library [55]. We also considered a class of benchmarks **sygus** corresponding to first-order formulations of the 71 single-invocation synthesis conjectures taken from the conditional linear integer track of the 2015 edition of the syntax-guided synthesis competition [1]. All benchmarks are in the SMT version 2 format. For comparisons with automated theorem provers, they were converted to the TPTP format by the SMTtoTPTP conversion tool [6]. We remark that all benchmarks consist purely of quantified formulas over linear arithmetic with very little, and in a majority of cases, no quantifier-free content. [6]

The results for the linear real and integer benchmarks are in Figures 9 and 10 respectively. Of the 7 benchmark classes, only one (the **scholl** class from LRA) had quantified formulas with nested quantification. The algorithm in Section 5 naturally extends to such formulas; a formal treatment of nested quantification is the subject of future work.

For LRA, we considered both the selection function from Figure 3, and its alternative from Figure 4, where the latter we refer to as CVC4 (a). For both LRA and LIA, the best configuration of CVC4 solves the most benchmarks overall (599 and 461 respectively), and did not give a conflicting response with any of the other solvers. Moreover, we note that CVC4 solves *every* benchmark that does not involve nested quantification,

[6] Details can be found at `http://cs.uiowa.edu/~ajreynol/InstLA`.

	keymaera (222)		**scholl** (371)		**tptp** (25)		**Total** (621)	
	#	time	#	time	#	time	#	time
CVC4(a)	**222**	2.0	**352**	2176.4	**25**	0.2	**599**	2178.6
CVC4	**222**	2.0	351	1074.0	**25**	0.2	598	1076.2
Z3	**222**	2.4	326	553.0	**25**	0.4	573	555.8
VampireZ3	220	51.2	57	393.2	**25**	2.3	302	446.7
Beagle	**222**	377.9	53	577.9	**25**	29.7	300	985.5
Vampire	218	57.8	43	31.1	**25**	1.3	286	90.1
Yices	**222**	0.9	–	0.0	**25**	0.01	247	1.0
ZenonArith	205	13.8	25	452.9	14	0.9	244	467.7
Princess	202	1136.2	0	0.0	**25**	67.4	227	1203.6

Fig. 9. Results for LRA benchmarks, showing times (in seconds) and benchmarks solved by each solver and configuration over 3 benchmark classes with a 300s timeout. Yices (version 2.4.1) does not support nested quantification, hence it was not applicable for the scholl class.

	psyco (189)		**tptp** (46)		**uauto** (155)		**sygus** (71)		**Total** (461)	
	#	time	#	time	#	time	#	time	#	time
CVC4	**189**	78.7	**46**	0.4	**155**	1.9	**71**	22.0	**461**	103.0
Z3	183	32.1	**46**	0.7	**155**	1.8	**71**	19.0	455	53.6
Beagle	28	900.0	**46**	48.4	153	343.6	57	617.7	284	1909.7
Princess	13	513.4	**46**	48.0	**155**	201.9	68	418.8	282	1182.1
VampireZ3	4	3.1	36	4.7	**155**	106.3	55	151.8	250	265.9
Vampire	6	196.0	36	2.0	**155**	378.0	46	262.8	243	838.7
ZenonArith	0	0.0	30	1.9	154	15.0	28	1374.6	212	1391.5

Fig. 10. Results for LIA benchmarks, showing times (in seconds) and benchmarks solved by each solver and configuration over 4 benchmark classes with a 300s timeout.

giving confirming evidence that our approach and implementation for solving linear arithmetic with one quantifier alternation is indeed sound and complete. Although we do not claim completeness for formulas with nested quantification, CVC4 solves more benchmarks (352) from the **scholl** class than any other solver.

The SMT solver Z3 (version 4.3.2), which uses the approach described in [11], solves the next most benchmarks overall, solving 573 and 455 total for the LRA and LIA sets respectively. A technique [18] in the SMT solver Yices (version 2.4.1) is able to solve all benchmarks from the **keymaera** and **tptp** classes of LRA, both does not handle quantified formulas in LIA or with nested quantification. We also considered the entrants of the first-order typed theorems division (TFA) of CASC 25, the most recent competition for automated theorem provers [56]. [7] For both benchmarks over LIA and LRA, the automated theorem provers trail the performance of CVC4 (and Z3) significantly. The best LRA automated theorem prover, VampireZ3, which uses a combination of a first-order theorem prover and an SMT solver [46], solves only 302 benchmarks, compared to 599 solved by CVC4 (a). The best LIA automated theorem

[7] We omit SPASS+T, which did not handle some classes of benchmarks due to restrictions on its input format, was comparable to the other automated theorem provers for the others. We show results for an updated version of Beagle (version 0.9.30).

prover, Beagle [7], solves 284 benchmarks, also notably less than the 461 solved by CVC4.

When comparing the best configuration of CVC4 to a combination of all other solvers, CVC4 solved 31 benchmarks that no other system solved, while in only 10 cases did another system solve a benchmark that CVC4 could not solve. In addition to solving the most benchmarks, CVC4 generally has small runtimes for the benchmarks it solves. When compared to the second best solver Z3, which takes 1129.8 seconds to solve 1028 benchmarks over all classes, CVC4 while solving 1060 benchmarks overall, solves its first 1028 benchmarks in a total of 159.0 seconds. Among the benchmarks solved by CVC4 and other systems, in only 12 cases did any system solve a benchmark at least 5 seconds faster than CVC4, while in 13 cases CVC4 solved a benchmark at least 5 seconds faster than all other systems.

Combining Linear Arithmetic with Uninterpreted Functions A prototype of the instantiation-based procedure from this paper was used by CVC4 in both the CASC J7 and CASC 25 competitions [56], which evaluated automated theorem provers on TPTP benchmarks involving combinations of arithmetic and free function symbols. CVC4 won the theorems (TFA) division of CASC J7 and finished 2^{nd} in CASC 25, behind VampireZ3. Additionally, CVC4 won the non-theorems (TFN) division of CASC 25. While treatment of uninterpreted functions is beyond the scope of this work, this shows the potential of the technique for use in general first-order automated theorem proving in the presence of background theories.

7 Conclusion

We have presented a class of instantiation-based procedures that are at the same time complete for quantified linear arithmetic and highly efficient in practice. Thanks to our framework we also obtain a simple and modular correctness argument for soundness and completeness on formulas with one quantifier alternation.

For future work, we would like to adapt the approach for quantified linear arithmetic with arbitrary quantifier alternations, and develop heuristics for avoiding worst case performance for quantified integer arithmetic involving large coefficients. We plan to develop selection functions for other theories, in particular, algebraic datatypes and fixed-width bitvectors, as well as for combinations of theories that admit quantifier elimination. A longer term goal of this work is to develop an approach that is effective in practice for quantified formulas involving both background theories and uninterpreted functions. We plan to investigate the use of the framework described in this paper as a component of such an approach.

Acknowledgements We would like to thank Peter Baumgartner for his help with converting the benchmarks used in the evaluation to the TPTP format.

References

1. R. Alur, R. Bodik, E. Dallal, D. Fisman, P. Garg, G. Juniwal, H. Kress-Gazit, P. Madhusudan, M. M. K. Martin, M. Raghothaman, S. Saha, S. A. Seshia, R. Singh, A. Solar-Lezama,

E. Torlak, and A. Udupa. Syntax-guided synthesis. To Appear in Marktoberdrof NATO proceedings, 2014.
2. R. Backofen. A complete axiomatization of a theory with feature and arity constraints. *Journal of Logic Programming*, 24:37–72, 1995.
3. K. Bansal, A. Reynolds, T. King, C. Barrett, and T. Wies. Deciding local theory extensions via e-matching. In *Computer Aided Verification (CAV)*. Springer, 2015.
4. C. Barrett, C. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Computer Aided Verification (CAV)*. Springer, 2011.
5. C. Barrett, A. Stump, and C. Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). `www.SMT-LIB.org`, 2010.
6. P. Baumgartner. Smttotptp a converter for theorem proving formats. In *CADE-25*, volume 9195 of *Lecture Notes in Computer Science*. Springer, 2015.
7. P. Baumgartner, J. Bax, and U. Waldmann. Beagle a hierarchic superposition theorem prover. In *Automated Deduction - CADE-25*, volume 9195 of *Lecture Notes in Computer Science*, pages 367–377. Springer, 2015.
8. L. Berman. The complexity of logical theories. *Theoretical Computer Science*, 11(1), 1980.
9. T. A. Beyene, S. Chaudhuri, C. Popeea, and A. Rybalchenko. A constraint-based approach to solving games on infinite graphs. In *POPL*, pages 221–234, 2014.
10. T. A. Beyene, C. Popeea, and A. Rybalchenko. Solving existentially quantified Horn clauses. In *CAV*, pages 869–882, 2013.
11. N. Bjørner. Linear quantifier elimination as an abstract decision procedure. In J. Giesl and R. Hähnle, editors, *IJCAR*, volume 6173 of *LNCS*, pages 316–330. Springer, 2010.
12. N. Bjørner and M. Janota. Playing with quantified satisfaction.
13. N. Bjørner, K. L. McMillan, and A. Rybalchenko. Program verification as satisfiability modulo theories. In *SMT@IJCAR*, pages 3–11, 2012.
14. H. Comon and C. Delor. Equational formulae with membership constraints. *Information and Computation*, 112(2):167–216, 1994.
15. D. C. Cooper. Theorem proving in arithmetic without multiplication. 1972.
16. L. M. de Moura and N. Bjørner. Efficient e-matching for SMT solvers. In F. Pfenning, editor, *CADE*, volume 4603 of *LNCS*, pages 183–198. Springer, 2007.
17. D. Detlefs, G. Nelson, and J. B. Saxe. Simplify: A theorem prover for program checking. Technical report, J. ACM, 2003.
18. B. Dutertre. Solving exists/forall problems with yices. In *Workshop on Satisfiability Modulo Theories*, 2015.
19. G. Fedyukovich, A. Gurfinkel, and N. Sharygina. Automated discovery of simulation between programs.
20. S. Feferman and R. L. Vaught. The first order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.
21. J. Ferrante and C. W. Rackoff. *The Computational Complexity of Logical Theories*, volume 718 of *Lecture Notes in Mathematics*. Springer, 1979.
22. H. Ganzinger and K. Korovin. New directions in instantiation-based theorem proving. In *Logic in Computer Science, 2003*. IEEE, 2003.
23. Y. Ge, C. Barrett, and C. Tinelli. Solving quantified verification conditions using satisfiability modulo theories. In *CADE*, volume 4603 of *LNCS*. Springer, 2007.
24. Y. Ge and L. de Moura. Complete instantiation for quantified formulas in satisfiability modulo theories. In *Proceedings of CAV'09*, volume 5643 of *LNCS*. Springer, 2009.
25. S. Grebenshchikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In *PLDI*, pages 405–416, 2012.
26. M. Heizmann, D. Dietsch, J. Leike, B. Musa, and A. Podelski. Ultimate automizer with array interpolation. In *TACAS*, 2015.

27. W. Hodges. *Model Theory*, volume 42 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1993.
28. S. Jacobs. Incremental instance generation in local reasoning. In *CAV '09*, pages 368–382, Berlin, Heidelberg, 2009. Springer-Verlag.
29. A. Komuravelli, A. Gurfinkel, and S. Chaki. SMT-based model checking for recursive programs. In *Computer Aided Verification*. Springer International Publishing, 2014.
30. D. Kozen. *Theory of Computation*. Springer, 2006.
31. V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Complete functional synthesis. In B. G. Zorn and A. Aiken, editors, *PLDI*, pages 316–329. ACM, 2010.
32. V. Kuncak and M. Rinard. Structural subtyping of non-recursive types is decidable. In *Eighteenth Annual IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, 2003.
33. R. Loos and V. Weispfenning. Applying linear quantifier elimination, 1993.
34. N. P. Lopes and J. Monteiro. Weakest precondition synthesis for compiler optimizations. In *VMCAI 2014*, pages 203–221, 2014.
35. M. J. Maher. Complete axiomatizations of the algebras of the finite, rational, and infinite trees. *IEEE Symposium on Logic in Computer Science*, 1988.
36. A. I. Mal'cev. *The Metamathematics of Algebraic Systems*, volume 66 of *Studies in Logic and The Foundations of Mathematics*. North-Holland, 1971.
37. D. Monniaux. Automatic modular abstractions for linear constraints. In *POPL 2009*, pages 140–151, 2009.
38. D. Monniaux. Quantifier elimination by lazy model enumeration. In T. Touili, B. Cook, and P. Jackson, editors, *CAV*, volume 6174 of *LNCS*, pages 585–599. Springer, 2010.
39. A. Mostowski. On direct products of theories. *Journal of Symbolic Logic*, 17(1), 1952.
40. T. Nipkow. Linear quantifier elimination. *Automated Reasoning*, pages 18–33, 2008.
41. A. Phan, N. Bjørner, and D. Monniaux. Anatomy of alternating quantifier satisfiability (work in progress). In *SMT 2012*, 2012.
42. A. Platzer, J.-D. Quesel, and P. Rümmer. Real world verification. In *Automated Deduction–CADE-22*, pages 485–501. Springer Berlin Heidelberg, 2009.
43. M. Presburger. über die vollständigkeit eines gewissen systems der aritmethik ganzer zahlen, in welchem die addition als einzige operation hervortritt. In *Comptes Rendus du premier Congrès des Mathématiciens des Pays slaves, Warsawa*, pages 92–101, 1929.
44. W. Pugh. The Omega test: a fast and practical integer programming algorithm for dependence analysis. In *ACM/IEEE conf. Supercomputing*, 1991.
45. C. R. Reddy and D. W. Loveland. Presburger arithmetic with bounded quantifier alternation. In *ACM STOC*, pages 320–325. ACM Press, 1978.
46. G. Reger, M. Suda, and A. Voronkov. Playing with avatar. In *Automated Deduction-CADE-25*, pages 399–415. Springer International Publishing, 2015.
47. A. Reynolds, M. Deters, V. Kuncak, C. W. Barrett, and C. Tinelli. Counterexample guided quantifier instantiation for synthesis in CVC4. In *CAV*. Springer, 2015.
48. A. Reynolds, T. King, and V. Kuncak. An instantiation-based approach for solving quantified linear arithmetic. *CoRR*, abs/1510.02642, 2015.
49. A. Reynolds, C. Tinelli, and L. D. Moura. Finding conflicting instances of quantified formulas in SMT. In *Formal Methods in Computer-Aided Design (FMCAD)*, 2014.
50. T. Rybina and A. Voronkov. A decision procedure for term algebras with queues. *ACM Transactions on Computational Logic (TOCL)*, 2(2):155–181, 2001.
51. C. Scholl, S. Disch, F. Pigorsch, and S. Kupferschmid. Using an smt solver and craig interpolation to detect and remove redundant linear constraints in representations of non-convex polyhedra. In *SMT*, pages 18–26. ACM, 2008.
52. T. Skolem. Untersuchungen über die Axiome des Klassenkalküls and über "Produktations- und Summationsprobleme", welche gewisse Klassen von Aussagen betreffen. Skrifter utgit av Vidnskapsselskapet i Kristiania, I. klasse, no. 3, Oslo, 1919.

53. T. Sturm and A. Tiwari. Verification and synthesis using real quantifier elimination. In *ISSAC 2011*, pages 329–336, 2011.
54. T. Sturm and V. Weispfenning. Quantifier elimination in term algebras: The case of finite languages. In *Computer Algebra in Scientific Computing (CASC)*, TUM Muenchen, 2002.
55. G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.
56. G. Sutcliffe. The 7th ijcar automated theorem proving system competition–casc-j7. *AI Communications*, pages 1–10, 2015.
57. A. Tarski. Arithmetical classes and types of algebraically closed and real-closed fields. *Bull. Amer. Math. Soc.*, 55(1), 1949.
58. A. Tarski. Arithmetical classes and types of boolean algebras. *Bull. Amer. Math. Soc.*, 55, 64, 1192, 1949.
59. R. Treinen. Feature trees over arbitrary structures. In P. Blackburn and M. de Rijke, editors, *Specifying Syntactic Structures*, chapter 7. CSLI Publications and FoLLI, 1997.
60. I. Walukiewicz. Monadic second-order logic on tree-like structures. *Theoretical Comput. Sci.*, 275(1–2):311–346, Mar. 2002.
61. V. Weispfenning. Complexity and uniformity of elimination in presburger arithmetic. In *ISSAC '97*, pages 48–53, New York, NY, USA, 1997. ACM.